

Article

A Dual-Axis Quantum-Resilience Framework for Autonomous Shipping

Simon Baradziej 

Department of Technology and Safety, UiT The Arctic University of Norway, Hansine Hansens veg 54, 9037 Tromsø, Norway; simon.baradziej@uit.no

Abstract

Quantum technology affects autonomous shipping along two paths that are usually studied separately. First, a cryptographically relevant quantum computer can break the public-key algorithms that authenticate maritime communications and adversaries can then harvest encrypted traffic to decrypt it later. Second, quantum gravimeters, magnetometers, and atomic clocks can help supply position and time when satellite navigation is jammed or spoofed. An autonomous vessel faces both of these situations at once with little human fallback. Based on the structured, reproducible review reported here, we found no framework that treats these two quantum-related axes as a single posture and scales it to vessel autonomy. This paper proposes such a framework. It is conceptual, not an empirically validated model. It defines a cryptographic axis and a resilient positioning, navigation, and timing axis, each read through anticipate, withstand, and recover and mapped to the four International Maritime Organization degrees of autonomy. Its central, testable claim is that the two axes are functionally non-substitutable, which we distinguish from statistical independence. A proposed autonomy-scaling principle, mediated by effective human-intervention capability, sets a non-compensatory required posture. A worked illustration, a real-vessel desk assessment, a validation programme, and a comparison with IMO, NIST, and ISO/IEC instruments are provided.

Keywords: autonomous shipping; maritime autonomous surface ships; post-quantum cryptography; resilient PNT; quantum sensing; positioning, navigation, and timing; resilience engineering; harvest-now-decrypt-later; GNSS interference



Academic Editors: Antonio M. Pascoal, Guoqing Zhang, Weidong Zhang and Jiqiang Li

Received: 5 July 2026

Revised: 20 July 2026

Accepted: 23 July 2026

Published: 25 July 2026

Copyright: © 2026 by the author.

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

Quantum technology affects autonomous shipping through two developments that are usually discussed by different communities. Cryptographers and cyber-risk specialists study the threat that a future quantum computer poses to public-key cryptography, along with the related concern that an adversary can record encrypted traffic today and decrypt it once such a machine exists. Physicists and navigation engineers study quantum sensing, meaning gravimeters, magnetometers, and atomic clocks that could help a vessel keep a trustworthy position and time when satellite navigation is jammed or spoofed. The first development is a threat to defend against. The second is a capability to acquire. The two are rarely considered together.

For a conventional ship with a full crew, considering them separately is reasonable. A bridge team can cross-check a doubtful position against radar, visual bearings, and a paper chart, and it can question an implausible instruction before acting on it. The crew is a general-purpose fallback that does not depend on which system failed. An autonomous

vessel does not have that fallback to the same degree, and the less human oversight it carries, the more it must manage both problems on its own. For autonomous shipping in particular, then, the two developments need to be considered together.

The regulatory context sharpens this point. In May 2026 the IMO adopted the first International Code of Safety for Maritime Autonomous Surface Ships (the MASS Code) at the 111th session of its Maritime Safety Committee (MSC 111, 13–22 May 2026) [1]. The Code applies to cargo ships and takes effect on 1 July 2026 on a non-mandatory basis, to support an experience-building phase; the IMO roadmap foresees adoption of a mandatory MASS Code by 1 July 2030, for entry into force on 1 January 2032 [1]. It organises autonomous operation into degrees, places strong emphasis on cybersecurity and on remote operations centres, and requires that the master retain overall responsibility for the ship at all times, even when not on board [1]. That responsibility is exactly what the two quantum axes can undermine.

This paper argues that the quantum resilience of an autonomous vessel depends on two capabilities that do not substitute for one another, and that the required level of each should be read against the vessel's degree of autonomy and, more precisely, against how much capable human intervention remains available. We develop the argument as a conceptual framework rather than a validated or calibrated model. We state this distinction plainly throughout, because the framework is proposed and has not been tested.

Three research questions organise the paper. RQ1 asks what the two quantum-related axes are, and how they differ as a threat and as a capability. RQ2 asks how the two axes can be combined into a single resilience posture and related to the degree of autonomy. RQ3 asks how the resulting claims can be stated so that they are testable. We answer RQ3 by casting the framework's central arguments as explicit propositions with operationally defined constructs (Section 4.4).

The contribution is a dual-axis quantum-resilience framework with five parts. It defines a cryptographic-resilience axis and a resilient-PNT axis, each read through anticipate, withstand, and recover. It represents a vessel's posture as a pair of ordinal levels on those axes. It relates the minimum required posture to the degree of autonomy through an effective-human-intervention-capability modifier. It states its arguments as testable propositions and gives a decision rule for the partially ordered space of postures. And it places the framework alongside existing IMO, NIST, and ISO/IEC instruments. Because the framework is conceptual, the paper also specifies the validation it would need.

The paper proceeds as follows. Section 2 sets out the review protocol, the evidence-weighting scheme, and the construction method. Section 3 reviews the cryptographic axis, the resilient-PNT axis, the IMO degrees of autonomy, and the resilience-engineering framing, assesses how well each suits the maritime autonomous context, and states the gap. Section 4 develops the framework, its propositions, the autonomy-scaling principle, the partial-order decision rule, and the treatment of dynamic autonomy. Section 5 gives a worked illustration, a cascading dual-failure scenario, a worked demonstration of the aggregation rules, and a desk assessment of a documented vessel. Section 6 draws out the implications and compares the framework with existing instruments. Section 7 states the limitations and prioritised threats to validity. Section 8 concludes with a plan for future work, and Appendix B gives an assessment checklist.

2. Materials and Methods

This is a conceptual contribution. Its result is a framework built from published evidence rather than from primary maritime operational data. The framework's claim to novelty depends on how complete that evidence is, and a reader should be able to judge and repeat how the evidence was gathered and used. This section therefore states the

review protocol, the scheme used to weigh evidence of different kinds, and the method that turned the reviewed evidence into the framework. The review is a structured narrative (scoping) review rather than a PRISMA-style systematic review. Its aim is coverage and synthesis across four adjacent studies, not an exhaustive enumeration of one.

2.1. Narrative Review Protocol

Sources were identified by searching Scopus, Web of Science Core Collection, IEEE Xplore, and the ACM Digital Library for peer-reviewed literature. These searches were supplemented by Google Scholar for grey literature, and by direct retrieval of primary standards and intergovernmental documents from NIST, the NSA, the IMO, the ITU, ISO/IEC, the IEC, and IACS. Searches were carried out between February and July 2026, and the final search date was 20 July 2026.

Representative Boolean strings were, for the cryptographic axis, (“post-quantum cryptography” OR “quantum-safe” OR “harvest now decrypt later” OR “crypto-agility”) AND (maritime OR ship* OR vessel OR AIS OR ECDIS); for the resilient-PNT axis, (“quantum sensing” OR “cold-atom” OR gravimeter OR magnetometer OR “atomic clock” OR “resilient PNT” OR “GNSS denial” OR “GNSS spoofing” OR “R-Mode” OR eLoran) AND (maritime OR ship* OR navigation); for autonomy, (“maritime autonomous surface ship” OR MASS OR “autonomous ship” OR “unmanned surface vehicle”) AND (autonomy OR “remote operation” OR “degree of autonomy”); and for resilience, (“resilience engineering” OR “cyber resilience” OR “anticipate withstand recover”) AND (maritime OR “safety-critical” OR system).

Records were included if they were peer-reviewed articles or conference papers, international standards, or authoritative intergovernmental, regulatory, or government technical reports; if they were written in English; and if they bore on at least one of the four themes in a way relevant to the maritime autonomous context. Foundational works were included without a date limit, for example Shor [2], while technology evidence was restricted to the period 2015 to 2026 to reflect current practice. Records were excluded when they lacked retrievable provenance or were promotional materials, with one exception: corporate announcements were kept only as evidence that a capability or trial exists (see Section 2.2). Titles and abstracts were screened for relevance to the maritime autonomous context, and the retained corpus is the set of sources cited in Sections 3–6. Because the search is narrative rather than systematic, the novelty statement in Section 3.6 is limited to the absence of a comparable framework in the searched corpus, and it is not a claim of universal absence.

To make the search reproducible, Table 1 reports the yield of a full-text search of the OpenAlex bibliographic database, run on 20 July 2026, for a representative query in each theme. These counts are broad and precede relevance screening; they establish that a search was run and can be repeated, not that every returned record is relevant. After removing duplicates and screening titles and abstracts for direct relevance to the maritime autonomous quantum context, the focused core corpus retained is the set of sources cited in Sections 3–6. The integration query, which combines all three concept clusters, returned records that mention the terms, but screening found none presenting a single non-substitutable dual-axis posture scaled to autonomy; this is the basis for the bounded novelty statement in Section 3.6.

Table 1. Reproducible search yield. Records identified by full-text search of the OpenAlex database on 20 July 2026 using the representative queries shown; counts are broad full-text matches, before relevance screening. After de-duplication and title/abstract screening for direct relevance to the maritime autonomous quantum context, a focused core corpus (the sources cited in Sections 3–6) was retained. No record returned by the integration query presented a single non-substitutable dual-axis posture scaled to autonomy, which bounds the novelty claim in Section 3.6.

Theme	Representative Query	Records Identified
Cryptographic axis	post-quantum cryptography maritime	708
Resilient–PNT axis	GNSS spoofing jamming maritime navigation	711
Degrees of autonomy	maritime autonomous surface ship	16,736
Resilience engineering	cyber resilience maritime autonomous ship	1929
Integration (this framework)	post-quantum cryptography quantum sensing autonomous ship	388

2.2. Evidence Weighting

Sources of very different authority bear on this topic, so evidence was weighed on an explicit hierarchy. Peer-reviewed literature ranked first, followed by international standards and intergovernmental or regulatory documents, then authoritative technical or government reports, and finally corporate announcements and press releases. Sources in the last tier were used only to establish that a capability or trial exists. They were not treated as performance evidence or as a basis for judging technological maturity. In particular, the Q-CTRL [3] sea-trial announcement is used here only as evidence that a maritime quantum-navigation trial took place. Its performance figures are attributed to the company and are not independently verified; claims about the maturity of civilian maritime quantum sensing rely instead on peer-reviewed sources [4,5]. This scheme is applied consistently in the text through hedged, source-attributed wording.

2.3. Framework-Construction Method

The framework was built by interpretive synthesis in six steps, each documented here. First, the reviewed exposures and capabilities were sorted into the smallest set of distinct, non-substitutable pairs, which produced two axes. Second, a common resilience lens of anticipate, withstand, and recover was taken from resilience engineering [6] and applied to each axis. Third, ordinal level descriptors (low, partial, high) were derived for each function from current practice, with anchors chosen to be observable. Fourth, the minimum required level on each axis was related to autonomy through the human-fallback argument, mediated by an effective-human-intervention-capability construct (Section 4.4). Fifth, the result was written as a partially ordered space of postures with an explicit decision rule (Section 4.6). Sixth, the construction was checked for internal consistency by tracing each modelling choice back to a reviewed source or a stated assumption, and its weak points were carried into the prioritised threats to validity in Section 7. The procedure is analytic. It produces a framework to be tested, not an empirical result.

Table 2 records this traceability: each modelling choice is set against the source or stated assumption that licences it, which makes the construction inspectable rather than asserted.

Table 2. Construction traceability: each modelling choice and the source or stated assumption that licences it (the internal-consistency check described in Section 2.3).

Modelling Choice	Basis (Source or Stated Assumption)
Two axes (cryptographic; resilient PNT)	Two non-substitutable quantum-related exposures and capabilities, drawn from the distinct literature (Sections 3.1 and 3.2; P1).
Functions: anticipate, withstand, recover	Adopted from cyber-resilience engineering (NIST, 2021 [6]), applied per axis (Section 3.4).
Three ordinal levels (low, partial, high)	Coarsest partition that separates the postures the argument needs; chosen for inter-rater tractability (Section 4).
Weakest-function aggregation (P5)	Chain-strength argument and safety-critical conservatism (Sections 4.4 and 7).
Non-compensatory Pareto rule (P6)	Avoids the single-index masking of single-axis exposure (Sections 4.6 and 5.4).
Cryptographic floor at partial for D1	HN DL and supply-chain exposures are present regardless of crew and are not human-recoverable (Section 4.5).
EHIC modifier on the autonomy baseline	The IMO degree is only a proxy for the availability of a capable human, which is what drives consequence (Sections 3.3 and 4.5).
Required mapping D1 to D4	Axis-specific human-fallback argument and the asymmetry of P4 (Section 4.5).

3. Background and Related Work

3.1. The Cryptographic Axis

Public-key cryptography underpins trust in digital systems. It lets two parties that have never met agree on a shared key, and it lets a recipient check that a message came from the claimed sender. Shor’s algorithm shows that a large, fault-tolerant quantum computer, a CRQC, would break the public-key schemes in common use [2]. The response is migration to PQC. In August 2024 NIST published three finalised standards: a module-lattice key-encapsulation mechanism [7], a module-lattice digital-signature scheme [8], and a stateless hash-based signature scheme [9]. National guidance places migration on the scale of the coming decade [10,11], and early maritime-specific analyses have appeared [12,13].

It helps to separate the security objectives that the threat touches, because their time-lines and consequences differ. The confidentiality of long-lived data is exposed through harvest-now-decrypt-later (HN DL): an adversary can record encrypted traffic today and read it once a CRQC exists, so data with a long secrecy lifetime is already at risk whenever the CRQC arrives. Authentication of ship-to-shore communications and the integrity of signed software and firmware are broken only once a CRQC can forge signatures, but the consequence of accepting a forged command or a tampered update is immediate and physical. Key and certificate lifecycle management, and recovery after compromise, cut across all of these objectives. The axis definition in Section 4.1 keeps them distinct.

Maritime systems sit awkwardly against this background, because several were not designed with authentication at all. The Automatic Identification System (AIS), which broadcasts a vessel’s identity and position, has no authentication by design and can be spoofed [14]. This is a classical vulnerability. It predates the quantum threat and is independent of it, and PQC does not by itself remedy it; that requires protocol and infrastructure modernisation. The IMO has begun exactly this modernisation: at the same session that adopted the MASS Code it introduced the VHF Data Exchange System (VDES), which adds authentication to ship tracking to reduce spoofing, for entry into force on 1 January 2028 [1]. That is a classical protocol fix, independent of the quantum threat, and it shows why the

framework separates classical vulnerabilities from the quantum-specific ones. Section 3.6 and Table A1 keep them apart.

3.2. The Resilient–PNT Axis

The second axis concerns keeping a trustworthy position, navigation, and timing solution when satellite navigation is denied. Denial is neither hypothetical nor rare, and it is not a quantum attack. Jamming overpowers the weak signals reaching a receiver, and spoofing replaces them with false ones [15]. The documented scale is large. One open-source investigation attributed nearly 10,000 GNSS spoofing events, affecting more than 1300 vessels across several sea areas, to a single state actor over about two years [16]. This grounds the axis in reported experience rather than in a hypothetical.

GNSS-independent PNT can be provided by conventional means and, in prospect, by quantum sensing. Conventional means include inertial navigation, terrestrial radionavigation such as ranging-mode (R-Mode) services on existing marine radiobeacons and eLoran, and interference monitoring. R-Mode, for instance, could cover a large share of the busiest sea areas if radiobeacons were upgraded to broadcast ranging signals [17]. Quantum sensing offers a further route: a cold-atom gravimeter or magnetometer matched against a stored map can yield position with no external signal, and an atomic clock can hold time through an outage [5,18]. Quantum sensing is therefore one enabling technology on this axis, not a requirement, and the axis is defined by resilient–PNT performance rather than by the presence of quantum hardware.

Two qualifications matter. First, gravity and magnetic fields are physical fields rather than broadcast signals, so they cannot be jammed or spoofed in the way an RF signal can. This does not make the axis immune to attack or error. The sensor, the reference maps, the calibration and processing software, and the sensor's own inputs can all be degraded, corrupted, or manipulated, and performance is sensitive to vessel motion and sea state. Second, the candidate technologies differ in maturity and error behaviour. Atomic clocks are mature. Strap-down inertial systems are mature but drift. Cold-atom gravimetric and magnetic sensing for civilian vessels is early and still largely defence-led [3,4]. Combining these technologies into one axis is justified by their shared function, which is to supply resilient PNT under GNSS denial, and not by any equivalence of readiness. The axis is scored on performance against the vessel's safe-navigation requirement (Section 4.2, Table A1).

3.3. Degrees of Autonomy

The IMO has organised autonomous operation into four degrees, developed through a regulatory scoping exercise [19] and carried into the MASS Code [1]. At Degree One a ship has automated processes and decision support, with seafarers on board operating the systems. At Degree Two it is controlled from another location, with seafarers on board. At Degree Three it is controlled from another location, with no seafarers on board. At Degree Four it is fully autonomous and decides and acts by itself. The degrees are an ordered abstraction of how much human oversight remains, and they are a useful organising structure. They are also only a rough proxy for whether a capable human is available at the moment of need, a limitation that Section 4.4 addresses with an explicit modifier rather than a caveat.

3.4. Resilience Engineering

Resilience engineering treats failure as expected and asks how a system behaves across the life of an adverse event, not only how well it resists one. A common formulation in the cyber domain distinguishes the abilities to anticipate, withstand, recover, and adapt [6]. This paper uses anticipate, withstand, and recover as the per-axis functions and treats

adaptation as cross-cutting, so that learning from recovery and incidents improves future anticipation on both axes. The choice is deliberate and has a consequence, discussed in Section 7: it captures single-event resilience well, but it represents adaptation to repeated or evolving threats less directly, and a system can recover from one event without adapting to the next.

3.5. Suitability of Each Body of Work for the Maritime Autonomous Context

Each of the four studies was developed for a setting that differs from an autonomous vessel, and it is worth saying how well each transfers. Post-quantum cryptography was developed mainly for enterprise and internet systems. Maritime differs in its long asset lives, its intermittent and bandwidth-limited connectivity, and its broadcast protocols that were never built for authentication, so migration priorities and feasibility differ from the enterprise case. Quantum sensing was developed in physics laboratories and defence programmes, and its transfer to civilian vessels is limited by size, weight, power, cost, and technology-readiness, so laboratory or defence performance cannot be assumed at sea. The IMO degrees are a regulatory abstraction that is useful for oversight but coarse as a measure of real human availability. Resilience engineering is a general systems discipline that has to be instantiated for the maritime case, which is what the per-axis functions and level descriptors attempt. Recognising these limits in fit is part of why the framework is offered as conceptual, and why its constructs are defined to be testable rather than assumed to hold.

3.6. Gap Statement

These four bodies of work are each well developed, but in the searched corpus they do not connect in the way an autonomous vessel needs. The cryptographic literature treats the quantum threat to authentication and confidentiality without reference to positioning. The quantum-sensing and resilient-PNT literature treats GNSS-independent position without reference to authentication. The autonomy literature classifies oversight without reference to either quantum axis. The resilience literature offers functions without a maritime instantiation that spans both axes. Within the corpus described in Section 2, we did not find a framework that treats the two quantum-related axes as a single, non-substitutable posture and scales that posture to the degree of vessel autonomy. A large control-theoretic literature does address resilience and disturbance rejection in autonomous ship control. Recent examples include finite-time disturbance-rejection dynamic positioning [20] and event-triggered distributed formation control of unmanned surface vehicles under communication delays and output constraints [21]. This work secures the control loop against physical disturbance and imperfect communication, on the assumption that its inputs are trustworthy. It is complementary to the present framework and operates downstream of it, because the framework concerns whether the inputs themselves, meaning authenticated messages and a trustworthy position and time, can be relied upon at all. The framework proposed here is meant to fill the integrative gap while connecting cleanly to that control layer (Section 6.3).

4. The Dual-Axis Quantum-Resilience Framework

The framework has a simple structure and a set of definitions that give it content. It uses two axes, cryptographic resilience and resilient PNT, each read through anticipate, withstand, and recover, and it represents a vessel's posture as the pair of ordinal levels it holds. Three design choices should be stated before the axes are defined. First, the levels are conceptual rather than numeric: the framework uses low, partial, and high rather than a score, because the evidence base does not yet support calibrated thresholds. Second, a

three-level scale is used rather than continuous or finer scoring, because three levels are the coarsest partition that still separates the postures the argument needs to tell apart, and because a coarse ordinal scale is more likely to be applied consistently by independent assessors than a fine or continuous one. The number of levels is itself a question for validation (Section 7). Third, the framework is proposed and not validated, so its content is written as required capability to be tested, not as a mature assessment standard.

4.1. Axis One: Cryptographic Resilience

The cryptographic axis measures how well a vessel and its shore support keep authentication and confidentiality trustworthy against the quantum threat. Because the threat touches objectives with different timelines, the axis is defined over four named subdimensions: confidentiality of long-lived data against HNDL; authentication of ship-to-shore communications; integrity of the software and firmware supply chain; and key and certificate lifecycle management with recovery after compromise. Each subdimension is read through anticipate, withstand, and recover. At low resilience the vessel relies on classical public-key cryptography only, holds no inventory, and has no migration or recovery plan. At partial resilience the vessel has begun a cryptographic inventory, identified its most exposed long-lived data and links, and deployed hybrid protection on those links as a transitional measure, but coverage is incomplete, and software and firmware signing usually lag. At high resilience the vessel keeps a complete inventory and mature crypto-agility, extends PQC or hybrid protection across authentication, long-lived data, and software and firmware signing, and has exercised distributed rekeying and certificate re-issuance across ship and shore. The step from partial to high is therefore the step from transitional hybrid deployment on the most exposed links to mature, exercised crypto-agility across all safety-relevant links. Because the four subdimensions differ in priority, the axis level is set by the weakest subdimension (Section 4.4, P5), and the subdimension profile is reported next to the axis level, so that a vessel well prepared for data-confidentiality migration but poorly prepared for firmware signing is visible rather than hidden inside an aggregate.

4.2. Axis Two: Resilient PNT

The resilient-PNT axis measures how well a vessel keeps a trustworthy position and time when satellite navigation is denied. It is read through anticipate—meaning recognising GNSS-denial risk and monitoring for interference and integrity faults; withstand—meaning holding a usable position and time through an outage; and recover—meaning reacquiring and cross-checking satellite signals before trusting them again. The level is judged on performance against the vessel's safe-navigation requirement, not on the technology used. High resilience requires a GNSS-independent position and time source that meets the route's accuracy and integrity needs and can carry the vessel through the longest plausible outage on that route. This can be met by capable conventional sensor fusion, such as inertial navigation aided by terrestrial radionavigation like R-Mode or eLoran with interference and integrity monitoring [17], by quantum sensing [5], or by a combination. Naming a technology is never enough to claim a level. The level is claimed with evidence that position, velocity, and timing accuracy, integrity or alert limits, and holdover duration meet the route's requirement; indicative anchors are given in Table A1. Where a required high level depends on a capability whose civilian maritime readiness is not yet established, the requirement is flagged as a capability gap rather than recorded as attainable today (Table A1, symbol legend).

4.3. Functional Non-Substitutability, Independence, and Common-Mode Failure

The framework uses two axes rather than one because the main countermeasures do not substitute for each other. PQC restores trust in authentication and confidentiality

but yields no position, and resilient-PNT capability yields a position but authenticates no message. This functional non-substitutability is the framework's central proposition (P1, Section 4.4). It must be kept separate from architectural or statistical independence, which the framework does not claim and which is generally false. The two axes can share resources such as trusted timing, electrical power, compute and software platforms, data buses and networks, shore communications, and maintenance and configuration governance, so common-mode and cross-axis failures are possible. A cyber compromise could corrupt the processing of quantum-sensor data, and the loss of a trusted time source could impair both navigation and the validation of credentials at the same time. Figure 1 therefore separates axis-specific measures from shared measures, and shows beneath both a layer of common-mode dependencies whose failure can degrade both axes together. The claim is deliberately narrow. The remedies are non-substitutable, while the coupling between the axes is acknowledged and built into the framework rather than assumed away.




Figure 1. Measures that serve the cryptographic axis only, the resilient-PNT axis only, and both, over a layer of common-mode dependencies whose failure can degrade both axes at once. A measure is axis-specific if it protects the function of only one axis, shared if it supports both, and common-mode if its failure can disable both. The placement is proposed and illustrative, and it reflects the author's classification rather than an empirical partition.

4.4. Propositions and Operational Definitions

To make the framework testable, its arguments are stated as propositions over defined constructs. Effective human-intervention capability (EHIC) is the extent to which a capable human can detect and correct a failure on either axis within the time the failure allows. It depends on the reachability and latency of a capable human, the staffing and workload of any remote operations centre, the onboard fallback available, the operational design domain and traffic complexity, and the length of time the vessel operates without human contact. Required resilience is the minimum axis level the framework prescribes for a given degree and EHIC. Failure consequence is the expected severity of an unrecovered failure on an axis. On these constructs, the framework advances six propositions.

P1, functional non-substitutability: no level on one axis reduces the required level on the other, because the main countermeasures address different functions. P2, autonomy-scaling monotonicity: as EHIC falls, the minimum required level on each axis does not decrease. P3, cryptographic floor: the cryptographic axis has a positive minimum of at least partial that does not fall to low even at maximum EHIC, because HNDL exposure and unverifiable software integrity cannot be recovered by human intervention. P4, asymmetry: the resilient-PNT minimum is more sensitive to EHIC than the cryptographic minimum, because a crew is a direct substitute for lost positioning but not for harvested data. P5, weakest-function within an axis: an axis level equals the lowest of its anticipate, withstand, and recover sublevels, and for the cryptographic axis the lowest of its subdimensions, because the resilience of a chain is bounded by its weakest link and a strong recovery cannot make up for absent anticipation. P6, non-compensatory across axes: a posture is adequate only if it meets or exceeds the required level on both axes, so a surplus on one axis does not compensate for a deficit on the other.

Of the six, P1, P5, and P6 follow by construction from the framework's definitions, while P2, P3, and P4 are the empirical conjectures that carry the substantive claim. The construction propositions can still be examined empirically: whether independent assessors apply the weakest-function rule consistently (P5), and whether the non-compensatory reading changes decisions relative to a single-index baseline (P6), are both measurable (Section 5.4). The empirical conjectures P2 to P4 can be tested through expert elicitation and, where incident data allow, against observed consequences. The functions, their cross-cutting adaptation, and the foundation that sets the required level are shown in Figure 2. Section 8 sets out the corresponding studies.




Figure 2. The resilience functions anticipate, withstand, and recover, applied separately to each axis, with adaptation shown as a cross-cutting feedback that carries recovery and incident experience back into future anticipation. The foundation combines the degree of autonomy with effective human-intervention capability (EHIC), which together set the required level on each function. Common-mode dependencies underlie all three functions (Figure 1).

To make EHIC usable rather than only qualitative, Table 3 gives an ordinal rubric. Each factor from the definition above is scored high, medium, or low, and the vessel's overall EHIC band is set by its weakest factor, mirroring the weakest-function rule. A

lower band raises the required posture on the axis where a reachable human can otherwise substitute, which is the resilient–PNT axis, while the cryptographic floor is unaffected (P3).

Table 3. Ordinal rubric for effective human-intervention capability (EHIC). The overall band is set by the weakest (lowest) factor, mirroring the weakest-function rule; a lower band raises the required posture on the axis where a reachable human can act. Thresholds are proposed and fixed by the validation programme (Section 8).

Factor	High EHIC	Medium EHIC	Low EHIC
Reachability and latency of a capable human	Continuous, low-latency contact	Intermittent or higher-latency contact	No contact for extended intervals
Remote-operations-centre staffing and workload	Dedicated, low workload per vessel	Shared, moderate workload	Minimal or none
Onboard human fallback	Crew aboard and able to act	Reduced crew or limited ability	No crew aboard
Operational design domain and traffic	Simple, low-density	Moderate	Complex, high-density
Endurance required without human contact	Minutes	Up to hours	Hours to whole voyage phases

For example, a Degree Three vessel in continuous low-latency contact with a dedicated, lightly loaded remote operations centre, operating in a simple domain, scores high EHIC on every factor; the framework then reads its resilient–PNT baseline one step below the uncrewed default, because a reachable operator can supply part of the anticipate and recover functions. The same vessel in a congested strait, with an overloaded centre and minute-scale intervals of no contact, scores low EHIC, and the full high-on-both requirement binds. EHIC thus turns the autonomy degree from a fixed label into an adjustable baseline, and it is the construct whose validation Section 8 prioritises.

4.5. The Autonomy-Scaling Principle and the EHIC Modifier

The scaling principle rests on the role of the human fallback. On a crewed ship the bridge team compensates for failures on both axes: an officer can question a suspicious message, or navigate by radar, visual bearings, and chart until a position is restored. As autonomy rises, that fallback thins. It is present at Degree One, degraded but real at Degree Two, only reachable ashore at Degree Three, and absent during any no-contact interval at Degree Four. From this, the level of resilience the vessel itself must hold rises as the fallback shrinks. The principle is about required resilience, not attained resilience, and it does not claim that higher autonomy attracts more attacks. It claims that the same failure has a worse expected consequence when fewer capable people are positioned to catch it.

Two refinements keep the principle honest and answer the objection that the IMO degree is only a rough proxy. First, the degree sets a baseline required posture that is then adjusted by EHIC, so a Degree Three vessel with a fast, well-staffed shore centre may carry a lower effective requirement than a poorly supported one at the same degree. The mapping is therefore not a rigid one-to-one relation, but a baseline modified by the quantity that actually drives consequence. Second, the human fallback is not a reliable detector of sophisticated deception. A crew can catch an obviously implausible position or instruction, but a well-crafted spoof or forged command that shows internally consistent false information on trusted displays may pass unchallenged. The fallback is a general-purpose remedy for overt failures, not a guarantee against covert manipulation, and this further supports the cryptographic floor.

With these refinements, the framework proposes a minimum posture for each degree, and it treats the axes asymmetrically because the fallback is axis-specific. The cryptographic axis keeps a floor that does not fall away at low autonomy. Even at Degree One the minimum is partial, because the HNDL and supply-chain parts of the threat are present whoever is on the bridge and cannot be reversed by the crew. The floor is partial rather than low because those non-recoverable exposures must be addressed. It is partial rather than high because a crewed vessel with reachable, well-resourced support can defer full crypto-agility while the CRQC timeline is uncertain, provided it holds an inventory and hybrid protection on its most exposed links. The resilient-PNT axis has no such floor and begins at low for Degree One, because a crew is a direct substitute for lost positioning. The two requirements converge as autonomy rises: partial on both at Degree Two, and high on both once no seafarers remain on board. Degrees Three and Four share the same required level, high on both axes, but they differ in where the resilience must reside and for how long it must hold unaided. At Degree Three a reachable human ashore may supply part of the anticipate and recover functions within a bounded latency. At Degree Four the vessel must sustain both axes on board through the longest planned no-contact interval, that is, as EHIC approaches zero. The value of scaling past Degree Three is thus not a higher level but a stricter continuity and endurance requirement, and this is what separates the two degrees in the framework and in Figure 3.




Figure 3. The minimum required posture by degree of autonomy on the ordinal posture grid. A vessel's posture is the pair of ordinal levels it holds on the cryptographic axis (horizontal) and the resilient-PNT axis (vertical). Markers D1 to D4 give the minimum posture the framework requires at IMO Degrees One to Four, and a posture is adequate for a degree only if it sits at or above the marker on both axes (Pareto dominance). The axes are ordinal and non-metric, and the connectors are discrete steps rather than a continuous trajectory. D3 and D4 share the same required level (high, high) but differ in where resilience resides: D3 permits a reachable human ashore, whereas D4 must sustain both axes on board through intervals with no human contact.

4.6. Posture as a Partial Order, and a Decision Rule for Incomparable Postures

Because each axis is an ordinal chain, the space of postures is their product and forms a partial order rather than a total order. Two postures can be incomparable: neither (high, low) nor (low, high) dominates the other. This is a feature, not a defect, and it gives a clean decision rule. A vessel complies at its degree and EHIC if, and only if, its posture Pareto-dominates the required posture, meaning it meets or exceeds the required level on both axes. If a vessel's posture is incomparable to the requirement, it is non-compliant on whichever axis falls short, and remediation is directed at that axis. This partial-order reading is what makes the dual-axis framework more informative than a single aggregate index. An index imposes a total order and can rate (high, low) as medium, and therefore adequate, which hides a complete exposure on the deficient axis. The non-compensatory rule (P6) refuses that masking and names the axis that must be fixed. Section 5.4 shows this contrast numerically.

4.7. Dynamic Autonomy and Phase-Dependent Posture

A real vessel may sit between degrees or change degree during a voyage, so a single static posture can mislead. The framework handles this by evaluating the required posture for each voyage phase and taking the binding requirement to be that of the most autonomous phase (lowest EHIC) the vessel will enter without reconfiguration. A vessel must therefore hold the highest required posture across its planned phases, unless it performs a verified reconfiguration between phases whose transitions have been tested and whose failure modes are covered. This rule resolves mixed and intermittent autonomy, such as a coastal vessel that is crewed on part of a route and uncrewed on another, by binding the requirement to the most demanding phase rather than to an average. It also makes dynamic switching an explicit design and assessment consideration rather than an unstated assumption.

5. Worked Illustration, Demonstration, and Desk Assessment

The illustration places two vessels on the framework. The values are illustrative and are labelled as such. They show how the framework reads a posture and how the requirement changes with autonomy, and they do not describe any real ship or measured capability.

5.1. Vessel A: A Remotely Operated Coastal Cargo Vessel

Vessel A is a short-sea cargo vessel operated from a remote operations centre, with a small crew on board for parts of its route, so it moves between Degree Two and Degree Three within a single voyage. By the phase rule (Section 4.7), the binding requirement is that of its uncrewed Degree Three phase, high on both axes, rather than an average over the voyage. Suppose its current posture is partial on the cryptographic axis and low on resilient PNT. On the cryptographic axis it has begun an inventory and placed hybrid protection on the main shore link, but software updates to onboard equipment are still validated with classical signatures, and recovery procedures are untested, so the firmware-signing subdimension holds the axis at partial. On the resilient-PNT axis it relies on satellite positioning, carries an inertial unit that drifts within minutes, and does not monitor for interference, which is low resilience. The gap is one step on the cryptographic axis and two steps on resilient PNT. Closing it means extending PQC or hybrid protection to software and firmware signing and exercising rekeying and re-issuance, as well as adding interference monitoring together with a GNSS-independent position and time source able to carry the vessel through a sustained outage on its route, with a tested reacquisition and cross-check procedure. The framework also makes clear that, when positioning is denied in a busy strait during an uncrewed phase, the right response may be a minimum-risk manoeuvre,

such as slowing, holding, or diverting to a safe state, rather than continued passage. High resilient-PNT resilience widens the set of situations in which safe continuation is justified, but it does not require it.

5.2. Vessel B: A Fully Autonomous Vessel

Vessel B operates at Degree Four and has intervals with no real-time human contact, during which no person can catch a forged command or an implausible position. Its required posture is high on both axes, with the resilience functions resident on board. Anticipation has to be automatic, with continuous monitoring of both the authenticity of inputs and the integrity of positioning. Withstanding has to be onboard, with PQC or hybrid protection of every safety-relevant input and a GNSS-independent position and time source that runs continuously. Recovery has to be self-directed, able to fall back to a safe state, reject a spoofed fix, and re-establish trust without waiting for a human. These are demanding functions, and the framework does not treat them as solved. Automatic detection and rejection require decision thresholds, quantified false-positive and false-negative rates, and verification evidence; specifying and testing those is part of what a Degree Four posture entails rather than something the framework grants. The illustration describes the posture the framework requires of such a vessel, and it makes no claim that one exists today.

5.3. A Cascading Dual-Failure Scenario

The value of the two-axis reading is clearest when the axes fail together. Consider Vessel B approaching a congested strait in a region with documented GNSS interference [16]. An adversary spoofs the position so that the false fix is internally consistent with a forged route command delivered over a classically protected link. Each failure then corroborates the other, and neither looks anomalous on its own. With resilience on only one axis, the vessel is still lost. Strong PQC alone would reject the forged command but leave the vessel with a plausible false position and no independent check. Strong resilient PNT alone would flag the position but still obey the forged command that redirects it. Only high resilience on both axes breaks the cascade: the independent position source contradicts the spoof, the signature check refuses the forged command, and the vessel holds a minimum-risk state until trust is restored. A single-axis or averaged treatment would misread this situation, whereas the dual-axis, non-compensatory reading captures it.

5.4. Worked Demonstration of the Aggregation and Compliance Rules

Because the framework is not yet calibrated, a worked demonstration is useful to show how a classification responds when individual criteria change, and to demonstrate the non-compensatory rule. The analysis is analytic and applies the framework's own aggregation rules to Vessel A's starting posture. It shows three things. First, the weakest-function rule (P5) governs the cryptographic axis: upgrading the main-link protection alone leaves the axis at partial, because firmware signing still lags, whereas upgrading firmware signing and exercising recovery is what raises it. Second, improvements confined to one axis do not by themselves produce compliance: raising resilient PNT from low to high while the cryptographic axis stays partial moves Vessel A to a posture of partial cryptographic and high PNT, which is still non-compliant with the Degree Three requirement of high on the cryptographic axis, whereas a single averaged index would have read the same change as a rise from low-to-medium to medium-to-high and might have declared it adequate. Third, the binding requirement is sensitive to the phase rule: if Vessel A were reclassified as Degree Two throughout, the requirement would fall to partial on both axes, and its only remaining gap would be on resilient PNT. Table 4 summarises these perturbations and their

effect on the posture and on compliance, and it shows why the framework's output should be read as the pair of axis levels together with the deficient axis, not as a single number.

Table 4. Worked demonstration: how Vessel A's classification responds to individual changes under the framework's aggregation and compliance rules (illustrative).

Perturbation from Vessel A Baseline	Crypto	Resilient-PNT	Posture	Compliant with D3 (High, High)?
Baseline (as assessed)	partial	low	(partial, low)	No. Short on both axes
Upgrade main-link protection only	partial	low	(partial, low)	No. Firmware-signing subdimension still caps crypto at partial (P5)
Upgrade firmware signing and exercise recovery	high	low	(high, low)	No. Resilient-PNT axis short
Raise resilient PNT to high only	partial	high	(partial, high)	No. Crypto short; a single averaged index would misread this as adequate (P6)
Upgrade both axes	high	high	(high, high)	Yes
Reclassify whole voyage as Degree 2	partial	low	(partial, low) vs. required (partial, partial)	No. Only the resilient-PNT gap remains

5.5. Desk Assessment of a Documented Vessel

To show the framework on a real architecture rather than an invented one, this subsection applies it to a documented vessel using only public information. The Yara Birkeland is a fully electric container vessel on a short coastal route in Norway; it is run with a remote operations centre and has completed a multi-year supervised-autonomy trial, with the stated aim of uncrewed operation, and its situational-awareness suite of radar, lidar, cameras, and AIS is supplied by its technology partner [22]. This is operator and vendor material, which sits in the lowest evidence tier (Section 2.2), so the levels below are inferences from public sources rather than verified assessments, and they would need operator data to confirm.

On autonomy and EHIC, the vessel today operates with supervision and a reachable centre—a higher EHIC than its uncrewed target; by the phase rule (Section 4.7) the binding requirement is the uncrewed phase—high on both axes. On the resilient-PNT axis, the public record shows a rich situational-awareness suite, but absolute position still rests on satellite navigation; there is no public evidence of a continuous GNSS-independent position-and-time source or of dedicated interference monitoring, so on public information the axis reads no higher than partial, short of the high its target autonomy would require. On the cryptographic axis, the public record emphasises cybersecurity and a cloud-linked centre but says nothing about post-quantum readiness, harvest-now-decrypt-later exposure, or software and firmware signing; here the framework's finding is that public information is insufficient to assign a level, which is itself useful, because it names precisely the evidence an assessor must request.

Two lessons follow. First, the framework applies to a genuine architecture and yields a specific, testable reading: a probable resilient-PNT gap and an undisclosed cryptographic posture. Second, its immediate value is to structure an assessment and surface the missing evidence, not to output a validated score, which is consistent with its status as proposed and turns the abstract illustration of Sections 5.1 and 5.2 into a concrete public example.

6. Implications and Synthesis

Because the framework is proposed and has no empirical findings, this section is written as implications and synthesis rather than as results.

6.1. Implications for Design, Operation, and Oversight

For design, the framework argues that quantum resilience should be specified as two requirements rather than one, and sized to the intended degree of autonomy and EHIC from the start rather than retrofitted. A design that satisfies one axis and neglects the other is not half resilient. It fails completely along the neglected axis. For operation, the framework frames a monitoring task that spans both axes, because a vessel needs to know in near real time both whether its inputs are authentic and whether its position can be trusted. The shared region in Figure 1 shows that some of this monitoring, along with governance, documentation, and training, can be organised once for both axes. Timing shows how careful the separation has to be. Precise time supports both positioning and some security functions, so a resilient clock helps on both axes at once. That is a common-mode dependency, not a collapse of the two axes, because a clock still authenticates no message and fixes no position by itself. For oversight, the framework speaks to class societies and flag administrations now that the MASS Code requires certification of remote operations centres and a master with genuine capability to intervene [1], because that capability is exactly what fails when either axis fails.

6.2. Relationship to Existing Frameworks: IMO, NIST, and ISO/IEC

The framework is a lens that composes existing instruments rather than a replacement for any of them. Its resilience functions come from NIST's cyber-resilience engineering [6], and its cryptographic axis operationalises the NIST PQC standards [7–9] and the associated migration guidance [10]. On the IMO side, the cryptographic axis maps onto existing maritime cyber-risk management expectations [23] and the MASS Code's intervention requirement [1], while the resilient-PNT axis currently sits across navigation safety and radio-spectrum protection rather than in any single instrument [24]. On the standards side, the cryptographic axis aligns with information-security management under ISO/IEC 27001 [25] and with the cyber-resilience requirements for ships and onboard systems in IACS UR E26 and E27 [26]. The dual-axis framework adds what none of these provides on its own: a single posture that couples the two quantum-related axes and scales the requirement to autonomy and EHIC. Table A1 and Figure 1 are where an assessor would connect the framework to these instruments in practice.

6.3. Relationship to Ship-Control Resilience and Anti-Interference Control

The framework sits upstream of the control layer. Control-theoretic work on autonomous ships secures the vessel's motion against physical disturbance and imperfect communication. Finite-time disturbance-rejection dynamic positioning under thruster saturation [20] and event-triggered distributed formation control under communication delays and output constraints [21] are representative. This work generally assumes the position, command, and timing inputs it consumes are trustworthy. The dual-axis framework addresses that assumption directly, because it concerns whether the inputs can be relied upon before the controller acts on them. The two are complementary, and the link between them is the path by which an input anomaly reaches the safety-decision and control layers. A spoofed position or a forged command that is not caught on the resilient-PNT or cryptographic axis becomes a corrupted input to an otherwise sound controller, which then tracks a false reference faithfully. Making that path explicit shows why input-trust resilience and control robustness have to be designed together.

6.4. Transferability

The structure of the framework is not specific to ships. Any autonomous safety-critical system that depends on authenticated communication and on satellite positioning, such as uncrewed aircraft and automated port equipment, faces the same two non-substitutable axes and the same rise in consequence as human oversight falls. The dual-axis reading and the EHIC-mediated scaling could therefore transfer, with the autonomy scale exchanged for the relevant one. This transfer is a direction for future work rather than a claim established here, and it is noted only to indicate the framework's likely scope.

7. Limitations and Threats to Validity

The framework is proposed and has not been validated. It is a synthesis of literature mapped to an autonomy scale and a resilience framing, its levels are conceptual descriptors rather than measured thresholds, and it has not been applied to a fleet, calibrated against incident data, or tested for whether independent assessors would place the same vessel at the same levels. This revision resolves in the manuscript itself several items that a conceptual contribution can settle without new data—namely, the review methodology (Section 2), the operational definitions and propositions (Section 4.4), the partial-order decision rule (Section 4.6), and the treatment of dynamic autonomy (Section 4.7). Only genuinely data-dependent items are left to validation.

The remaining threats to validity are given in order of importance. First and most important, the construct validity of EHIC and of the level thresholds is untested; if these constructs are not interpreted consistently, the framework's outputs will not be reproducible. Second, functional non-substitutability is defended, but the common-mode and cross-axis dependencies of Section 4.3 mean architectural independence does not hold, and shared-resource failures could dominate real behaviour. Third, the autonomy-scaling monotonicity (P2) is plausible but unchecked against incident data, so it is still possible that consequence does not rise smoothly as EHIC falls. Fourth, the three-level scale may prove too coarse to separate vessels that differ in practice, or too fine to be applied consistently. Fifth, the weakest-function, non-compensatory aggregation (P5, P6) is conservative by design, justified by the chain-strength argument and by the safety-critical setting, but the choice to allow no weighting or compensation is itself a hypothesis to be tested. Sixth, treating adaptation as cross-cutting rather than as a fourth per-axis function captures single-event resilience but represents adaptation to repeated or evolving threats less directly. Seventh, the evidence base for civilian maritime quantum sensing is thin and shows existence rather than maturity, so the resilient-PNT axis should be read as required capability, not present availability.

Alternative structures were considered, and they are compared here rather than only listed. A single aggregate index was rejected because it imposes a total order and hides single-axis exposure, as Section 5.4 shows. A four-function scheme that promotes adaptation to a per-axis function is a workable extension and would represent evolving threats better, at the cost of a larger matrix that is harder to apply; it is the most likely refinement if validation shows the three-function scheme misses repeated-threat behaviour. Continuous scoring was set aside for now because the evidence base does not support calibrated thresholds, though it becomes attractive once instrument development (Section 8) yields measurable indicators. Additional autonomy modifiers beyond EHIC were considered, but EHIC was adopted as the single modifier that most directly captures the quantity that drives consequence, which is human availability at the moment of need.

8. Conclusions

Quantum technology meets autonomous shipping along two axes that the literature usually treat apart. One is the threat that quantum computing poses to maritime authentication and confidentiality, answered by migration to post-quantum standards. The other is the capability that resilient PNT offers when satellite navigation is denied, a capability for which quantum sensing is one emerging enabler. This paper has argued that an autonomous vessel is exposed on both axes at once, that neither capability substitutes for the other, and that the level of resilience a vessel needs on each should rise as effective human intervention falls with increasing autonomy. The dual-axis quantum-resilience framework expresses this as a posture read on two axes, each defined through anticipate, withstand, and recover, mapped to the IMO degrees of autonomy through an EHIC modifier, stated as testable propositions, and equipped with a partial-order decision rule.

The framework is conceptual, and its value will be settled by validation. A concrete programme follows. Expert elicitation would convene a panel spanning naval architecture, maritime cyber security, PNT and navigation engineering, remote-operations practice, and class and flag surveying on the order of twelve to eighteen experts, and would run a three-round Delphi to test whether the axes, levels, functions, and the EHIC modifier are recognised as the right ones, and to fix the level thresholds. The objective, stated as an objective rather than an assumed property, is that independent assessors converge on the same classification. Case application to real or planned autonomous vessels, including those entering the MASS Code experience-building phase, would be selected to span the degrees and a range of operational design domains, and would test whether the framework captures the postures that arise in practice and whether the scaling principle holds. Inter-rater reliability would be measured with a chance-corrected statistic such as Krippendorff's alpha or a weighted Cohen's kappa, with a target of at least 0.7 and explicit rules for missing and conflicting indicators. Instrument development would turn the conceptual levels into a scored assessment with evidence requirements for each level, starting with the cryptographic axis, which is the most ready for measurable scoring, and defining performance indicators for each axis. Until those steps are taken, the contribution is a way of seeing the problem as a whole: two quantum-related axes, one posture, scaled to how much the vessel must look after itself.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

AIS	Automatic Identification System
CRQC	Cryptographically relevant quantum computer
EHIC	Effective human-intervention capability
GNSS	Global Navigation Satellite System
HNDL	Harvest-now-decrypt-later
IACS	International Association of Classification Societies
IMO	International Maritime Organization
MASS	Maritime Autonomous Surface Ship(s)
PNT	Positioning, navigation, and timing
PQC	Post-quantum cryptography
ROC	Remote operations centre
TRL	Technology-readiness level

Appendix A. Dual-Axis Resilience Assessment Matrix

Table A1 sets out the matrix. The two axes appear as row groups, each divided into the anticipate, withstand, and recover functions. The columns give the conceptual level descriptors (low, partial, high) with operational anchors, and the expectation related to the degree of autonomy. The expectation column restates, at per-function granularity, the required postures shown in Figure 3. The cells contain proposed content, read together with Sections 4.1–4.7 and subject to the same status: proposed and not yet validated. Where a required high level depends on a capability whose civilian maritime readiness is not yet established, the cell carries a capability-gap flag.

Table A1. Dual-axis resilience assessment matrix (proposed; not empirically validated). Thresholds are indicative and are to be fixed by the validation programme in Section 8. “Safe-navigation requirement” means the route-specific accuracy, integrity or alert limit, and outage endurance the vessel must meet.

Axis	Function	Low	Partial	High	Expectation by Degree of Autonomy
Cryptographic axis	Anticipate	No cryptographic inventory; harvest-now-decrypt-later (HNDL) exposure of long-lived data unrecognised.	Inventory begun; most exposed long-lived data and links identified; migration planned.	Complete inventory across all safety-relevant links; crypto-agility maintained; HNDL exposure tracked and reviewed.	D1 ● D2 ● D3 ● D4 ●
Cryptographic axis	Withstand	Classical public-key cryptography only.	Transitional hybrid protection on the most exposed links (e.g., the primary ship-to-shore link); coverage incomplete, software and firmware signing typically still classical.	PQC or hybrid protection across authentication, long-lived data, and software and firmware signing (all safety-relevant links); mature crypto-agility.	D1 ● D2 ● D3 ● D4 ●
Cryptographic axis	Recover	No rekeying or re-issuance capability.	Recovery procedures defined but not exercised.	Distributed rekeying and certificate re-issuance exercised across ship and shore within a defined recovery-time objective.	D1 ● D2 ● D3 ● D4 ●
Resilient–PNT axis	Anticipate	GNSS trusted by default; no interference or integrity monitoring.	Route GNSS-denial risk assessed; basic interference monitoring.	Continuous interference and integrity monitoring of positioning inputs against defined alert limits.	D1 ○ D2 ● D3 ● D4 ●
Resilient–PNT axis	Withstand	GNSS only; no independent position or time source.	Short-duration backup (e.g., inertial), where ‘short-duration’ is defined per route as the time to reach a safe state; not sufficient for a sustained outage (an outage longer than that time).	GNSS-independent position and time source meeting the route’s safe-navigation requirement and sustaining it through the maximum plausible outage on that route (see note a).	D1 ○ D2 ● D3 ● D4 ●▲

Table A1. Cont.

Axis	Function	Low	Partial	High	Expectation by Degree of Autonomy
Resilient–PNT axis	Recover	No reacquisition or cross-check of satellite signals.	Manual reacquisition after an outage, performed by the onboard crew or, where none is aboard, by the remote operations centre; may be unavailable when no capable human is reachable in time.	Automatic reacquisition with cross-check against the independent source, using defined decision thresholds, before trust is restored.	D1 ○ D2 ● D3 ● D4 ●

Note a: High resilient–PNT resilience may be met by conventional sensor fusion (inertial navigation with R-Mode or eLoran) or by quantum sensing; the civilian-vessel readiness of the quantum options is not yet established (Section 3.2), so the D4 expectation carries a capability-gap flag. Symbol legend. The level the framework proposes as required at each IMO degree of autonomy (D1 to D4): ● high required; ● partial acceptable; ○ low acceptable; ▲ capability gap or civilian-vessel readiness not yet confirmed. The symbols denote the required level only. They do not encode a vessel's attained level, the associated uncertainty, or the present commercial availability of the enabling technology, each of which is recorded separately in an assessment.

Appendix B. Assessment Checklist

This checklist turns the framework into a step-by-step assessment for a single vessel. It needs no data beyond what an operator and its suppliers can provide, and it produces a posture, a compliance verdict, and a prioritised gap list. The desk assessment in Section 5.5 shows that recording what evidence is missing is part of the output.

Step 1. Establish the autonomy profile. Record the IMO degree in each planned voyage phase, and identify the lowest-EHIC phase the vessel enters without reconfiguration (Section 4.7).

Step 2. Score EHIC. Using the rubric in Table 3, score each factor high, medium, or low for that phase; the overall band is the weakest factor.

Step 3. Read the required posture. From the degree and the EHIC band, take the minimum required posture from Figure 3 and Table A1; the cryptographic floor is at least partial regardless of EHIC.

Step 4. Score the cryptographic axis. For each of the four subdimensions (Section 4.1), score anticipate, withstand, and recover against the Table A1 anchors; the axis level is the weakest sublevel across all subdimensions (P5). Record the subdimension profile, not only the axis level.

Step 5. Score the resilient–PNT axis. Score anticipate, withstand, and recover against the anchors, judging withstand and recover on measured performance against the route's safe-navigation requirement, not on the technology named (Section 4.2).

Step 6. Apply the compliance rule. The vessel complies only if its posture meets or exceeds the required level on both axes (Pareto dominance, P6); if not, name the deficient axis or axes.

Step 7. Record gaps and evidence needs. List the upgrades that would close each gap, and flag any required level that depends on a capability whose readiness is not yet established (the capability-gap flag in the Table A1 legend).

References

1. International Maritime Organization. *IMO Adopts First Global Code for Autonomous Ships (MASS Code)*; Press Briefing; IMO: London, UK, 2026. Available online: <https://www.imo.org/en/mediacentre/pressbriefings/pages/imo-adopts-mass-code.aspx> (accessed on 20 July 2026).
2. Shor, P.W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **1997**, *26*, 1484–1509. [CrossRef]

3. Q-CTRL. Q-CTRL's New Maritime Quantum Navigation Solution Successfully Undergoes First Defense Trials at Sea; Press Release; Q-CTRL: Sydney, Australia, 2025. Available online: <https://q-ctrl.com/blog/q-ctrls-new-maritime-quantum-navigation-solution-successfully-undergoes-first-defense-trials-at-sea> (accessed on 20 July 2026).
4. Bongs, K.; Holynski, M.; Vovrosh, J.; Bouyer, P.; Condon, G.; Rasel, E.; Schubert, C.; Schleich, W.P.; Roura, A. Taking atom interferometric quantum sensors from the laboratory to real-world applications. *Nat. Rev. Phys.* **2019**, *1*, 731–739. [\[CrossRef\]](#)
5. Li, H.; Zaminpardaz, S.; Kealy, A.; Greentree, A.D.; Rubinov, E.; Gibson, B.; Choy, S. Quantum sensors for enhanced positioning and navigation: A comprehensive review. *GPS Solut.* **2026**, *30*, 62. [\[CrossRef\]](#)
6. National Institute of Standards and Technology. *Developing Cyber-Resilient Systems: A Systems Security Engineering Approach*; NIST Special Publication 800-160, Rev. 1; U.S. Department of Commerce: Gaithersburg, MD, USA, 2021; Volume 2. [\[CrossRef\]](#)
7. FIPS 203; Module-Lattice-Based Key-Encapsulation Mechanism Standard. U.S. Department of Commerce: Gaithersburg, MD, USA, 2024. [\[CrossRef\]](#)
8. FIPS 204; Module-Lattice-Based Digital Signature Standard. U.S. Department of Commerce: Gaithersburg, MD, USA, 2024. [\[CrossRef\]](#)
9. FIPS 205; Stateless Hash-Based Digital Signature Standard. U.S. Department of Commerce: Gaithersburg, MD, USA, 2024. [\[CrossRef\]](#)
10. Cybersecurity and Infrastructure Security Agency; National Security Agency; National Institute of Standards and Technology. *Quantum-Readiness: Migration to Post-Quantum Cryptography*; Cybersecurity Information Sheet; CISA: Washington, DC, USA, 2023. Available online: <https://media.defense.gov/2023/Aug/21/2003284212/-1/-1/0/CSI-QUANTUM-READINESS.PDF> (accessed on 20 July 2026).
11. National Security Agency. *Announcing the Commercial National Security Algorithm Suite 2.0*; Cybersecurity Advisory PP-22-1338, Version 1.0; NSA: Fort Meade, MD, USA, 2022. Available online: https://media.defense.gov/2025/May/30/2003728741/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS.PDF (accessed on 20 July 2026).
12. Berger, D.; Lye, A.; Maidl, A.; Stoppe, J.; Windhorst, A. Post-quantum cryptography for maritime systems. In *Proceedings of the 2025 IEEE International Conference on Cyber Security and Resilience (CSR)*, Uppsala, Sweden, 4–6 August 2025; IEEE: New York, NY, USA, 2025. [\[CrossRef\]](#)
13. Tabish, N.; Chaur-Luh, T. Maritime autonomous surface ships: A review of cybersecurity challenges, countermeasures, and future perspectives. *IEEE Access* **2024**, *12*, 17114–17136. [\[CrossRef\]](#)
14. Balduzzi, M.; Pasta, A.; Wilhoit, K. A security evaluation of AIS automated identification system. In *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC'14)*, New Orleans, LA, USA, 8–12 December 2014; Association for Computing Machinery: New York, NY, USA, 2014; pp. 436–445. [\[CrossRef\]](#)
15. Psiaki, M.L.; Humphreys, T.E. GNSS spoofing and detection. *Proc. IEEE* **2016**, *104*, 1258–1270. [\[CrossRef\]](#)
16. Center for Advanced Defense Studies (C4ADS). *Above Us Only Stars: Exposing GPS Spoofing in Russia and Syria*; C4ADS: Washington, DC, USA, 2019. Available online: <https://c4ads.org/reports/above-us-only-stars/> (accessed on 20 July 2026).
17. Koch, P.; Gewies, S. Worldwide availability of maritime medium-frequency radio infrastructure for R-Mode-supported navigation. *J. Mar. Sci. Eng.* **2020**, *8*, 209. [\[CrossRef\]](#)
18. Wright, M.J.; Anastassiou, L.; Mishra, C.; Davies, J.M.; Phillips, A.M.; Maskell, S.; Ralph, J.F. Cold atom inertial sensors for navigation applications. *Front. Phys.* **2022**, *10*, 994459. [\[CrossRef\]](#)
19. International Maritime Organization. *Outcome of the Regulatory Scoping Exercise for the Use of Maritime Autonomous Surface Ships (MASS)*; MSC.1/Circ.1638; IMO: London, UK, 2021.
20. Gao, X.; Long, Y.; Li, T.; Hu, X.; Chen, C.L.P.; Sun, F. Optimal fuzzy output feedback control for dynamic positioning of vessels with finite-time disturbance rejection under thruster saturations. *IEEE Trans. Fuzzy Syst.* **2023**, *31*, 3447–3458. [\[CrossRef\]](#)
21. Ning, J.; Jin, L.; Chen, C.L.P.; Tong, S. Event-triggered distributed formation cooperative tracking control of unmanned surface vehicles with communication delays and output constraints. *IEEE Trans. Veh. Technol.* **2026**, *in press*. [\[CrossRef\]](#)
22. Yara International. Yara Birkeland. Available online: <https://www.yara.com/news-and-media/media-library/press-kits/yara-birkeland-press-kit/> (accessed on 20 July 2026).
23. International Maritime Organization. *Guidelines on Maritime Cyber Risk Management*; MSC-FAL.1/Circ.3/Rev.3; IMO: London, UK, 2025.
24. International Civil Aviation Organization; International Maritime Organization; International Telecommunication Union. *Joint Statement on Protecting the Radio Navigation Satellite Service (RNSS) from Harmful Interference*; ITU: Geneva, Switzerland, 2025. Available online: <https://www.itu.int/en/mediacentre/Documents/2025/ICAO-IMO-ITU-Joint-Statement.pdf> (accessed on 20 July 2026).

25. ISO/IEC 27001:2022; Information Security, Cybersecurity and Privacy Protection. Information Security Management Systems. Requirements. ISO: Geneva, Switzerland, 2022.
26. International Association of Classification Societies (IACS). *Cyber Resilience (UR E26) and Cyber Resilience of On-Board Systems and Equipment (UR E27)*; IACS: London, UK, 2022.

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.