Corroboration Is Not Authentication: Bounding the Receiver-Side Integrity Assurance Available to Unauthenticatable Cooperative Datalinks in Aviation and Maritime Transport

Jherrod Thomas

Independent Researcher — The Lion of Functional Safety

jherrodthomas.com · September 2026

Vol. 1, No. 10 · September 1, 2026 · 9 pages · 27 references

***Abstract*—**Two cybersecurity advisories published in August 2026 describe safety-of-navigation systems whose remediation field reads "no mitigation available." CISA ICSA-26-219-01 assigns five CVE identifiers to the absence of message authentication in Controller–Pilot Data Link Communications over ATN-B1, a property of an internationally standardized datalink rather than a defect in any vendor's product. CISA ICSA-26-237-07 assigns two CVE identifiers to a Class B Automatic Identification System transponder whose Maritime Mobile Service Identity can be rewritten over the vessel's own network, on a product discontinued in October 2020 with no software maintenance channel. In both cases the safety argument does not collapse; it silently migrates. Integrity that the protocol cannot supply is tacitly reallocated to receiver-side corroboration — voice readback, the controller's traffic picture, radar association, operator judgment — and that reallocation is nowhere written down as a requirement, given a coverage target, or verified. We formalize corroboration as an allocatable safety mechanism and derive its ceiling. Five numbered relations convert a corroboration ledger into an effective coverage figure, apply a common-cause coupling factor, restrict the coverage to the subset of channels an adaptive adversary cannot shape, impose a time-to-irreversible-commitment budget analogous to the fault-tolerant time interval, and return a residual undetected-falsification rate. The central result is a bound rather than an estimate: against an adversary who can author the message, only channels outside the adversary's write authority contribute assurance, and channels whose verdict arrives after commitment contribute nothing at all. We instantiate the framework on both advisories, produce corroboration ledgers for an uplinked altitude clearance and for own-ship identity broadcast, and derive six requirements anchored in DO-326A/ED-202A, ED-120, IEC 63154, and IACS UR E26/E27.

Index Terms*—Controller–Pilot Data Link Communications, Automatic Identification System, DO-326A, ED-202A, IEC 63154, Cooperative Datalink, Message Authentication, Plausibility Monitoring, Safety–Security Co-Assurance, Residual Risk*

I. Introduction

A safety analyst confronted with an unauthenticated datalink has three honest options: change the protocol, stop using it, or write down what is holding the integrity claim up instead. The literature has spent two decades on the first option and produced a body of workable cryptographic retrofits. Operations has never seriously entertained the second. The third — auditing the assurance that receiver-side plausibility checking actually supplies — is the one that neither community has formalized, and it is the one that every deployed fleet is silently relying on right now.

The occasion for this paper is a pair of advisories published three weeks apart. On August 7, 2026, CISA published ICSA-26-219-01 against Controller–Pilot Data Link Communications (CPDLC) over ATN-B1, all versions affected, five CVE identifiers, mitigation available: none [1]. The lead finding, CVE-2025-71409, is categorized CWE-306, missing authentication for a critical function, and its text is an engineering statement rather than a security one: the absence of authentication for VHF Data Link messages allows rogue ground stations to inject CPDLC messages leading to unexpected or misleading clearances. Four companion CVEs cover unauthenticated link-control frames that terminate sessions individually and by broadcast, and injected status messages accepted as legitimate [1]. On August 25, 2026, CISA published ICSA-26-237-07 against the FURUNO FA-50 Class B Automatic Identification System (AIS) transponder: CVE-2026-59769, use of hard-coded credentials, CVSS v3.1 base score 9.1 with confidentiality impact None and integrity impact High; and CVE-2026-67578, missing authentication for critical function, base score 7.5 [2]. Production of the affected unit ended in October 2020, the vendor states that software updates will no longer be provided, and the recommended residual control is that the vessel be properly locked and managed [2].

Neither advisory describes a system that stopped working. Both describe a system whose integrity claim was never written down. The aviation advisory is careful to say that the findings do not constitute an unsafe aircraft condition but can degrade operational safety margins by increasing workload, delaying safety-critical instructions, and reducing situational awareness [1]. That is a defensible statement, and it is defensible for a

specific reason that the advisory does not state: the flight crew and the controller are expected to catch the injected clearance. The maritime advisory's residual control is a door key, which is defensible for the same unstated reason: somebody is expected to notice that the vessel's identity is wrong.

Both defenses are corroboration arguments. Neither is an authentication argument, and the two are not interchangeable. Authentication is a property of the message: it either carries a verifiable binding to an authorized originator or it does not, and the property is independent of who is watching. Corroboration is a property of the *deployment*: it depends on which other channels are available at the moment of receipt, how much discriminating power each of them has against a falsified assertion, whether the adversary can shape any of them, and whether their collective verdict arrives before the consumer of the message does something irreversible. Corroboration can be excellent. It can also be zero, and it can be zero without anyone noticing, because nothing in the applicable process standards requires it to be enumerated.

Our claim is that this is a tractable engineering problem with a defensible bound, and that the bound is frequently much lower than practitioners assume. The dominant reason is not that individual plausibility checks are weak. It is that most of the checks a receiver performs are computed over fields the adversary authored. Format validity, sequence numbering, checksum agreement, session-state consistency, and — in the maritime case — the identity field itself are all inside the attacker's write authority. They discriminate perfectly against noise and not at all against an adversary who has read the specification. Against an adaptive adversary, effective coverage is set by the channels outside that write authority, and the count of such channels is often one or two, sometimes zero.

The contributions of this paper are as follows.

- A **corroboration ledger**, an enumeration artifact that lists for each asserted attribute of a message the independent channels available to the receiver, each channel's discrimination coverage, and — the field practitioners omit — whether the adversary can shape the channel's input.
- Five **numbered relations** that convert the ledger into a naive coverage figure, a common-cause-corrected figure, an adversary-restricted bound we call the *corroboration assurance bound*, a time-to-irreversible-commitment constraint analogous to the fault-tolerant time interval of automotive practice, and a residual undetected-falsification rate.
- A **structural result**: shapeable channels contribute nothing to the bound, so the assurance available to an unauthenticatable link is determined entirely by the receiver's access to sensing or judgment the adversary cannot write into, and by whether that access closes in time.
- A **worked application** to both August 2026 advisories, producing ledgers for an uplinked altitude clearance and for own-ship AIS identity, and six derived requirements traceable to DO-326A/ED-202A, ED-120, IEC 63154, and IACS UR E26/E27.

Section II reviews the background and identifies the gap. Section III develops the framework. Section IV applies it. Section V discusses limitations and where the standards lens stops. Section VI concludes.

II. Background and Related Work

A. Cooperative Links Specified Without Identity

Three broadcast or near-broadcast links carry safety-relevant assertions in transport today, and all three were specified without message authentication. Automatic Dependent Surveillance–Broadcast (ADS-B) has the most thoroughly documented security posture: Strohmeier, Lenders, and Martinovic's survey enumerated the injection, modification, deletion, and jamming attack classes and established that the protocol offers no defense at the message layer [3]; Wu, Shang, and Guo's later survey organizes the decade of proposed countermeasures that followed [4]; and Manesh and Kaabouch quantified the resulting risk posture across attack classes [5]. Sampigethaya and Poovendran had already framed the general problem as one of aviation cyber-physical systems in which the cyber layer's assumptions and the physical layer's consequences are analyzed by different communities using incompatible vocabularies [6].

CPDLC's position is narrower and worse. Gurtov, Polishchuk, and Wernberg documented the absence of authentication and encryption in the deployed datalink and set out the threat classes in 2018 [7]. Smailes and colleagues demonstrated practical attacks with commodity software-defined radio at CPSS 2021, including a man-in-the-middle takeover of an aircraft's session on an attacker-chosen frequency after which arbitrary CPDLC messages could be delivered without alerting the legitimate controller [8]. A 2025 review in the *Journal of Transportation Security* consolidates the attack taxonomy and concludes that cryptographic message authentication is the missing control [9]. Most recently, a full-stack analysis presented at USENIX Security 2026 built an SDR-based CPDLC ground station and demonstrated uplink injection and large-scale denial of service against certifiable avionics hardware in a laboratory environment [10]. ICSA-26-219-01 is the point at which this research record acquired CVE identifiers [1].

AIS occupies the maritime equivalent. Balduzzi, Pasta, and Wilhoit's ACSAC 2014 evaluation demonstrated ship spoofing, false distress-beacon injection, and vessel-identity manipulation against both the protocol and the online tracking providers built on it [11]. A 2026 survey synthesizes the intervening decade and characterizes the root causes as open VHF broadcast, self-reported data, absent authentication and replay protection, and GNSS dependence [12]. The FA-50 advisory adds a variant the survey literature under-treats: the falsification need not occur over the air at all, because the transponder's identity field is writable from the vessel's own Ethernet segment [2].

B. Cryptographic Retrofits and the Fleet-Lifetime Gap

Proposals to add authentication to these links are neither scarce nor unserious. Goudossis and Katsikas developed an identity-based public-key scheme for AIS that preserves the message format [13]; Kessler demonstrated a protected-AIS capability providing message integrity, timing integrity, and sender authentication [14]; and the VHF Data Exchange System (VDES) offers a forward path with authentication built in, now entering the IMO regulatory framework with carriage provisions expected to take effect in 2028 [15]. On the aviation side, the countermeasure surveys catalogue PKI-based and lightweight-authentication designs for both ADS-B and CPDLC [4], [9].

The gap this paper addresses is not that these designs are unsound. It is that none of them discharges the risk carried by the *installed base* within the interval that matters. VDES carriage provisions expected in 2028 do not change what a Class B transponder discontinued in October 2020 transmits tomorrow [2], [15]. A PKI for CPDLC does not change what ATN-B1 accepts on an active frequency this afternoon [1]. The advisories' "no mitigation available" is, read precisely, a statement that the protocol-layer options are unavailable *to the fleet in service*. Everything that protects that fleet in the interim is receiver-side, and that is exactly the layer no standard requires anyone to enumerate.

C. Corroboration as a Detection Strategy

The technical means of receiver-side corroboration are well developed, and it is worth noting that they were developed as *detection research* rather than as *assurance allocation*. Strohmeier, Lenders, and Martinovic showed that wide-area multilateration can validate ADS-B position claims against independently measured time-difference-of-arrival, because an adversary can control what an aircraft reports but not where a transmission physically originates [16]. The same logical structure appears in navigation: Psiaki and Humphreys's treatment of GNSS spoofing and detection organizes defenses by which physical property of the signal the spoofer cannot simultaneously satisfy [17], and receiver autonomous integrity monitoring generalizes the idea to cross-checking redundant measurements. In the maritime domain, the AIS anomaly-detection literature reviewed by Ribeiro, Paes, and de Oliveira applies kinematic and behavioral models to reported tracks [18], and machine-learning intrusion detection for ADS-B continues to develop [19].

Every one of these techniques is a corroborating channel in the sense used here. What the literature does not supply is the allocation step: a statement of how much assurance a given deployment's channels collectively provide, expressed in a form a safety case can consume, with the adversary's ability to shape individual channels made explicit.

D. Safety and Security Co-Assurance

The co-engineering literature has correctly identified that safety analysis and security analysis produce incompatible artifacts and must be reconciled. Kriaa and colleagues surveyed the combined approaches available to industrial control systems [20]; Lisova, Šljivo, and Čaušević's systematic review characterized thirty-three co-analysis methods by lifecycle stage, standards association, and validation status [21]; and Kavallieratos, Katsikas, and Gkioulos surveyed the cyber-physical co-engineering landscape more broadly [22]. Johnson and Kelly's assurance framework for independent co-assurance, and its later Bayesian-network instantiation against IEC 61508 and Common Criteria, address the structural question of how two assurance arguments with different epistemic bases can be combined without one silently absorbing the other [23], [24].

This body of work supplies the frame within which our contribution sits, and it stops short of the specific quantity we need. Co-analysis methods tell an analyst that a security-caused failure must be represented in the safety argument. They do not tell the analyst how much credit a plausibility monitor may claim against an adversary who authored the message the monitor is inspecting. That number is what a residual-risk statement requires and what the two August 2026 advisories implicitly assert without computing.

E. The Gap in One Sentence

No current standard requires the enumeration of the corroborating channels on which an unauthenticatable link's integrity claim depends, and consequently no current safety case states the coverage, the adversary-shapeability, or the latency of the mechanism that is actually holding the argument up.

III. Approach: The Corroboration Assurance Bound

A. Objects and Definitions

Let a *message* m carry an *asserted attribute* a — an altitude clearance, a reported position, a vessel identity — to a *consumer* that will take an action conditioned on a. Let $S = \{1, \ldots, N\}$ index the *corroborating channels* available to the receiver: sources of evidence, other than the message itself, that bear on whether a is true.

Definition 1 (Discrimination coverage). Channel i has *discrimination coverage* c_i in the unit interval against a falsification of a if, conditioned on a being falsified, channel i produces a verdict inconsistent with a with probability c_i.

Definition 2 (Shapeability). Channel i is *shapeable* by an adversary A if A can influence the channel's input so as to make its verdict consistent with the falsified a. Write $U \subseteq S$ for the *unshapeable* subset.

Definition 3 (Time to irreversible commitment). For a consumer action conditioned on a, t_{commit} is the interval from message acceptance to the first moment at which the action's effects can no longer be withdrawn without harm.

Definition 2 is the load-bearing one and it is the field practitioners routinely omit. A checksum has near-unit coverage against random corruption and zero coverage against an adversary, because the adversary computes it. A sequence number, a session identifier, a Current Data Authority designation, and a self-reported identity field are all in the same category. They are not weak checks; they are checks against the wrong initiating event. Fig. 1 draws the resulting architecture, with the shapeable channels distinguished from the unshapeable ones.

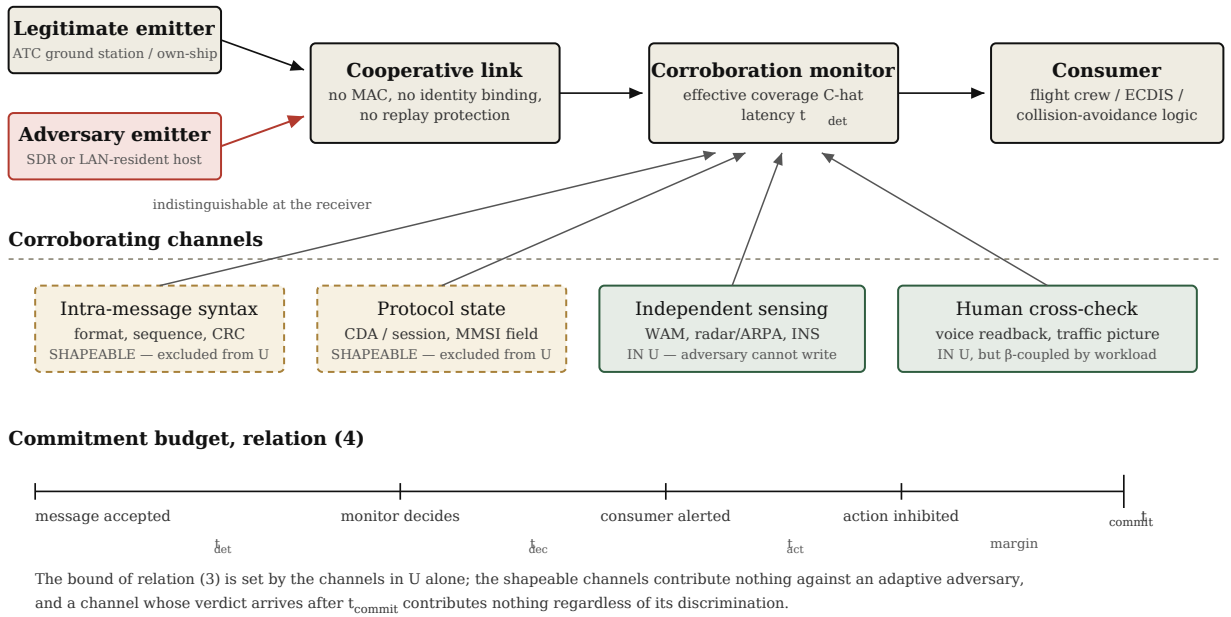

Figure 1. Receiver-side corroboration for an unauthenticatable cooperative link. The legitimate and adversary emitters are indistinguishable at the link layer, which is what the missing-authentication CWE means. Channels computed over fields the adversary authored are shapeable and are excluded from the set U used in relation (3); only independent sensing and human cross-check remain, and the latter is coupled to the former through operator workload, which is what the factor beta in relations (2) and (3) carries. The lower timeline shows the constraint of relation (4): a verdict arriving after the commitment point supplies no assurance whatever its coverage.

B. Naive and Common-Cause-Corrected Coverage

If the channels were mutually independent, the probability that at least one produces an inconsistent verdict would be the standard series-detection expression

$$C_{\text{naive}} = 1 - \prod_{i \in S} (1 - c_i) \tag{1}$$

Relation (1) is the figure practitioners reach for and it is optimistic in two distinct ways. The first is ordinary common-cause dependence, familiar from hardware safety analysis: a single condition — an operator's saturated attention, a shared time reference, a common GNSS source, a single display surface — can defeat several channels at once. We import the beta-factor treatment used in dependent-failure analysis and write the corrected coverage as

$$\hat{C} = (1 - \beta)\,[1 - \prod_{i \in S} (1 - c_i)] + \beta \cdot \max_{i \in S} c_i \tag{2}$$

where β in the unit interval is the fraction of falsification events in which the channels fail as a group rather than individually. The second term states the conservative position that under a common-cause condition the ensemble is worth no more than its single best member.

C. The Bound Under an Adaptive Adversary

The second optimism is the one this paper is about. Relations (1) and (2) both sum over S, which tacitly asserts that every enumerated channel is an independent witness. Against an adversary who authored m, the shapeable channels are not witnesses at all: their verdicts are functions of the adversary's own output. Restricting the product to U yields the *corroboration assurance bound*

$$\hat{C}_{\text{adv}} = (1 - \beta)\,[1 - \prod_{i \in U} (1 - c_i)] + \beta \cdot \max_{i \in U} c_i, \quad U \subseteq S \tag{3}$$

Relation (3) is austere and its austerity is the result. It contains no security-specific tuning term. It says only that an adversary who can write a channel's input cannot be detected by that channel, and it follows immediately that $\hat{C}_{\text{adv}} \leq \hat{C} \leq C_{\text{naive}}$, with the gap between the first and third quantities equal to whatever credit the analysis was taking from the attacker's own arithmetic. When U is empty, relation (3) returns zero and the correct statement about the deployment is that it has no integrity mechanism for a at all — not a weak one.

D. The Commitment Budget

Coverage without timeliness is not a mechanism. Automotive practice formalizes this as the fault-tolerant time interval and its decomposition into fault-detection and fault-reaction budgets; the same decomposition applies here, with the difference that the reaction is frequently performed by a human. Let t_{det} be the interval from message acceptance to the monitor's verdict, t_{dec} the interval from verdict to the consumer's awareness, and t_{act} the interval required to inhibit or reverse the action. The channel contributes assurance only if

$$t_{\text{det}} + t_{\text{dec}} + t_{\text{act}} \leq t_{\text{commit}} \tag{4}$$

Any channel violating (4) is removed from U before relation (3) is evaluated. This is where a great many real corroboration

arguments quietly fail. Post-hoc analytics on an AIS track corpus, an ANSP-side correlation of anomalous multi-aircraft logoffs, or a maintenance-time audit of transponder configuration are all valuable and all satisfy (4) for essentially no consumer action of interest.

E. Residual Rate and the Six-Step Procedure

Let Λ_{att} be the assumed rate of falsification attempts against a in the operational context, and P_{act} the conditional probability that an undetected falsified assertion produces the hazardous consumer action. The residual is

$$\Lambda_{\text{res}} = \Lambda_{\text{att}} \cdot (1 - \hat{C}_{\text{adv}}) \cdot P_{\text{act}} \tag{5}$$

Relation (5) is included for completeness and carries the heaviest caveat, developed in Section V: Λ_{att} is not a failure rate and should not be treated as one. Its function is to keep the residual on an axis the safety case already uses, not to forecast attacks.

The procedure is six steps.

1. **Fix the assertion and the consumer action.** Name the attribute a and the specific action conditioned on it. Different actions on the same message have different t_{commit} and different P_{act}.
2. **Enumerate channels.** List every source of evidence other than m that bears on a. Include human ones explicitly.
3. **Classify shapeability.** For each channel, ask whether the adversary who authored m can also set the channel's input. Every field carried inside m is shapeable by construction.
4. **Assign coverage and timing.** Assign c_i and evaluate (4). Channels failing (4) are struck.
5. **Evaluate (2) and (3).** Report both. The distance between them is the assurance the deployment was claiming from the adversary's own arithmetic.
6. **Derive requirements.** Where $\hat{C}_{\text{adv}}$ is insufficient for the hazard's classification, the requirement is either a new unshapeable channel, a coverage or latency improvement on an existing one, or an extension of t_{commit} by procedural inhibition.

Table I applies steps 1 through 3 at the link level to the three cooperative links discussed in Section II.

IV. WORKED EXAMPLE

A. An Uplinked Altitude Clearance over ATN-B1

Step 1. The assertion a is "the sector controller has cleared this aircraft to flight level X." The consumer action is the flight crew commanding the level change. Following the advisory's threat description for CVE-2025-71409 [1], the adversary is an SDR-equipped rogue ground station within radio horizon, transmitting a well-formed ATN-B1 uplink on the active VDL Mode 2 frequency. The commitment point is not the crew's acceptance of the message; it is the moment the aircraft's trajectory has diverged far enough that recovery interacts with adjacent traffic. In a high-density sector with reduced vertical separation this is on the order of tens of seconds after the level change is commanded, which sets a demanding t_{commit}.

Steps 2 and 3. Table II is the corroboration ledger. Five channels are available and three of them are shapeable, which is the finding.

Steps 4 and 5. With all five channels credited and no shapeability correction, relation (1) returns $C_{\text{naive}} \approx 0.998$, which is the number an uncritical review would report and is the reason the failure condition looks Minor. Restricting to $U = \{4, 5\}$ and taking $\beta = 0.2$ to represent the workload condition that degrades crew plausibility checking and voice-confirmation discipline simultaneously — a coupling explicitly recognized in the advisory's own language about increased workload [1] — relation (3) returns

$$\hat{C}_{\text{adv}} = 0.8 \times [1 - (0.45)(0.05)] + 0.2 \times 0.95 = 0.972$$

The apparent conclusion is comfortable and it is wrong, because channel 5 fails relation (4) under precisely the conditions that make the attack worth mounting. A crew that commands the level change on receipt and confirms by voice afterward has a voice verdict arriving after t_{commit}. Striking channel 5 leaves $U = \{4\}$ and relation (3) collapses to $\hat{C}_{\text{adv}} = 0.55$. The deployment's integrity mechanism for an uplinked clearance is one human plausibility check with slightly better than even odds, and the difference between 0.998 and 0.55 is the entire content of this paper.

The aggregate case is worse still. CVE-2025-71411 describes broadcast control frames disconnecting multiple aircraft simultaneously [1]. ED-120's availability allocation for CPDLC is sound for random failure because a proven voice reversion exists; its load-bearing assumption is that outages are uncorrelated with operational demand. An adversary supplies the correlation. Channel 5's coverage is a function of controller availability, and a sector-wide simultaneous reversion to a single voice frequency during a peak window is exactly the condition under which that availability collapses. In the framework's terms, the denial-of-service CVEs are not a separate hazard from the injection CVE; they are a mechanism for driving β toward unity against the injection CVE.

B. Own-Ship Identity on a Class B AIS Transponder

Step 1. The assertion a is "this vessel's MMSI is M." There are two distinct consumers and they must be analyzed separately, which is the structural feature this case contributes.

For the *own-ship* consumer — the master, who must know what identity the vessel is broadcasting — Table III's ledger is empty. A Class B AIS transponder is a broadcast-only transmitter. There is no loopback, no receive-own-transmission path, no configuration-change annunciation, and no alarm sentence covering a static-data write [2]. U is empty, and relation (3) returns $\hat{C}_{\text{adv}} = 0$ exactly. This is not a weak mechanism. It is the absence of one, and it is why the advisory's residual control has to be physical access management: with no detection channel at all, prevention is the only remaining category.

Table 1. Cooperative links carrying safety-relevant assertions in transport, with the asserted attribute, the authentication status at the message layer, the principal unshapeable corroborating channel available to a receiver, and the public-record advisory or literature anchor. The pattern the table makes visible is that the unshapeable channel is in every case a physical measurement or a human process belonging to a different duty-holder than the emitter.

Link	Asserted attribute	Message-layer authentication	Principal unshapeable channel	Anchor
ADS-B (1090ES)	Own aircraft position, identity, intent	None	Wide-area multilateration; primary radar	[3], [16]
CPDLC over ATN-B1	Controller clearance or instruction	None; CRC against noise only	Voice channel readback; controller traffic picture	[1], [7]–[10]
AIS Class B (own-ship transmit)	Own vessel identity and static data	None; identity is a configuration field	None onboard — transmitter is broadcast-only	[2], [11], [12]
AIS Class A/B (receive)	Other vessel identity, position, kinematics	None	Radar/ARPA target association; visual bearing	[12], [18]

Table 2. Corroboration ledger for an injected CPDLC uplink clearance under CVE-2025-71409 [1]. Coverage figures are structured engineering judgment for illustration, not measurements; Section V develops the consequences of that status. The three shapeable channels are the ones a receiver-side design review would most naturally credit, and they are exactly the ones relation (3) discards.

#	Channel	Basis	Shapeable	c_i	Meets (4)
1	Message syntax and CRC	Field carried in m	**Yes**	0 vs. adversary	Yes
2	Message reference number and closed-loop protocol state	Field carried in m	**Yes**	0 vs. adversary	Yes
3	Current Data Authority / session binding	Field carried in m; session takeover demonstrated in [8], [10]	**Yes**	0 vs. adversary	Yes
4	Crew plausibility against flight plan, terrain, and current request state	Airborne state, not adversary-writable	No	0.55	Yes
5	Voice confirmation with the sector controller	Separate channel and duty-holder	No	0.95	Only if initiated before commanding the change

Step 6 and the regulatory asymmetry. The FA-50's type-approval basis names IEC 62287-1 and ITU-R M.1371, which treat MMSI as a static installation parameter set once at commissioning against a ship station licence [2]. That requirement was satisfied at the human interface — the unit has no front-panel MMSI editor — and unsatisfied at the network interface, which is the gap the two CVEs occupy. IEC 63154:2021, the cybersecurity standard for shipborne navigational and radiocommunication equipment, addresses exactly this class of defect, and was published after production of the affected unit ended in October 2020 [2]. IACS UR E26 and E27 became mandatory for ships contracted for construction on or after 1 July 2024, and the USCG cybersecurity rule at 33 CFR Part 101 Subpart F took effect on 16 July 2025 for vessels already carrying a security plan. A Class B transponder on a workboat is in scope of none of them.

The framework makes the consequence precise rather than rhetorical. AIS is a shared safety service whose achieved $\hat{C}_{\text{adv}}$ at any receiver is set by the least-instrumented participant transmitting into the same VHF data link, because the receiving vessel's corroboration channels are the only ones that exist and they do not distinguish a Class A hull inside four regulatory regimes from a Class B hull inside none.

C. Derived Requirements

Table IV states the six requirements the two ledgers produce.

R4 deserves a comment because it is nearly free and would have changed the FA-50 finding's character entirely. The advisory's integrity impact is High on both CVEs and its confidentiality impact is None on both [2], which is the CVSS vector's crude way of saying that this is an authenticity finding on a device whose only output is data. A device that cannot prevent the write can still refuse to make it silent, and an annunciated write moves $\hat{C}_{\text{adv}}$ for the own-ship consumer from exactly zero to something bounded below by the operator's attention.

V. DISCUSSION

Coverage figures are judgment, not measurement. Every c_i in Tables II and III is structured engineering judgment. This is a real limitation and it is worth being precise about its scope. The framework's principal result — that shapeable channels contribute nothing and that late channels contribute nothing —

Table 3. Corroboration ledger for own-ship AIS identity falsification under CVE-2026-59769 and CVE-2026-67578 [2], separated by consumer. The own-ship row set is empty by architecture: the transponder does not receive its own broadcast, so no channel exists at any coverage or any latency. The receiving-vessel channels exist but are conditioned on a target being within radar or visual range and on an operator performing association.

Consumer	Channel	Shapeable	c_i	Meets (4)	$\hat{C}_{\text{adv}}$
Own ship	*None available*	—	—	—	**0**
Receiving vessel	Radar/ARPA target association with AIS track	No	0.7 in range	Yes	0.5 in range; 0 beyond
Receiving vessel	Visual identification of hull or name	No	0.3, daylight and close range	Yes	
Receiving vessel	Kinematic plausibility of reported track [18]	Partly — adversary controls reported kinematics	0.2	Yes	

is *structural* and does not depend on the numbers: it holds for any assignment of c_i. What depends on the numbers is only the magnitude of $\hat{C}_{\text{adv}}$ once the ledger has been pruned. An analyst who distrusts our figures should still perform steps 1 through 3, because the pruning is where the finding lives. Human-factors literature could supply defensible priors for the crew and controller channels, and multilateration and radar-association coverage are empirically measurable; neither has been done for this purpose.

The coupling factor is the weakest quantitative element. Beta-factor estimation is contested even in hardware dependent-failure analysis, where the coupling mechanisms are physical and enumerable. Here the dominant coupling is operator workload, which the adversary can *induce* — a property with no analogue in random-failure analysis. Relation (2) treats β as a parameter; the CPDLC case in Section IV-A shows it behaving as a decision variable available to the attacker. A more faithful formulation would make β a function of the adversary's denial-of-service capability, and we regard this as the framework's most important open problem.

Attempt rates are not failure rates. Relation (5) borrows the form of a hazard-rate expression and should not be read as one. There is no defensible base rate for Λ_{att}, and the advisories are explicit that no public exploitation has been reported [1], [2]. We include (5) because safety cases are organized around rates and a residual expressed in no units at all does not enter the conversation; we do not claim it is predictive. The conservative use is to fix Λ_{att} at one and read (5) as a conditional statement about what happens when an attempt occurs.

Where the standards lens stops. The most consequential limitation is jurisdictional rather than technical. Both cases place the decisive corroborating channel under a different duty-holder than the equipment. The CPDLC crew and controller channels belong to the operator and the ANSP; the airborne security process specification governs the aircraft. The AIS receiving-vessel channels belong to *other vessels entirely*, over which the falsified hull's owner has no authority and no visibility. A requirement written under DO-326A can compel an aircraft manufacturer to enumerate channels; it cannot compel the ANSP to keep a controller free enough to answer a voice call. This is the structural reason both advisories close with no mitigation available: the mitigation exists, it is just not in anyone's scope.

Threats to validity. The two cases were selected because they are recent, public, and share the no-mitigation property, which is a deliberate but non-random selection. The framework has not been applied to a case where the corroboration argument turns out to be strong, and such a case would be a more informative test than either presented here. The relations assume a single adversary with a fixed write authority; a coordinated adversary able to compromise a corroborating sensor moves that channel from U to $S \setminus \{\} U$, which the framework represents but does not help the analyst anticipate.

VI. Conclusion and Future Work

When a cooperative datalink cannot be given message authentication within the service life of the fleet that depends on it, the integrity claim does not disappear. It moves to the receiver, where it is carried by plausibility checks, independent sensing, and human judgment that no process standard currently requires anyone to enumerate. This paper has argued that the resulting assurance is boundable, and that the bound is set by two conditions with no free parameters: an adversary who authors the message cannot be detected by any check computed over that message, and a verdict arriving after the consumer has committed supplies nothing. Applied to CPDLC over ATN-B1, the bound reduces a nominally near-certain ensemble of five checks to a single human plausibility judgment. Applied to own-ship identity on a Class B AIS transponder, it returns exactly zero, which correctly reclassifies the finding from a weak-control problem to a no-control problem and explains why the vendor's residual measure is a physical one.

Three directions follow. First, the coverage figures should be replaced with measurements: crew and controller discrimination against injected clearances is amenable to simulator study, and radar-to-AIS association coverage is amenable to analysis of existing traffic corpora. Second, the coupling factor should be reformulated as an adversary-controlled quantity, which would let a single model express the observation that the denial-of-service CVEs in ICSA-26-219-01 are an enabling condition for the injection CVE rather than a separate finding. Third, the corroboration ledger should be evaluated as a candidate work product for the airworthiness security process and for

Table 4. Requirements derived from the two ledgers, with the relation that motivates each and the standards clause family under which it would be written. R1 through R3 attach to the airborne and ANSP ends of a CPDLC deployment; R4 through R6 attach to AIS equipment design and onboard integration.

ID	Requirement	Motivating relation	Clause family
R1	The security risk assessment shall enumerate, for each uplink message class capable of changing the vertical or lateral profile, the corroborating channels available to the crew, and shall record for each channel whether it is computed over fields carried in the message.	(3), Definition 2	DO-326A / ED-202A threat condition identification
R2	Flight crew procedure shall require voice confirmation to precede the commanded change for any uplinked clearance that alters the vertical profile and was not preceded by a corresponding downlink request.	(4)	Operational procedure; ED-120 RCP assumptions
R3	The ANSP shall detect and annunciate correlated multi-aircraft CPDLC session termination within a sector, and shall treat such an event as a precondition raising the coupling factor applied to crew and controller cross-check.	(2), (3)	DO-355 / ED-204 continuing airworthiness security; ANSP safety management
R4	Shipborne equipment that broadcasts an identity shall annunciate any change to that identity's stored value at the operator interface, and shall persist the change with a timestamp in a log readable without vendor tooling.	(3) with U empty	IEC 63154 general requirements
R5	The vessel's onboard network shall place identity-bearing transmitting equipment in a zone separate from general-purpose hosts and shared display equipment, with the conduit between them enumerated and access-controlled.	Definition 2	IACS UR E26 / E27; IEC 62443-3-2 zone and conduit
R6	Where equipment is out of software support, the residual risk statement shall record the achieved corroboration coverage for each safety-relevant assertion the equipment makes, and shall state explicitly where that coverage is zero.	(3), (5)	Residual risk disclosure; ISM cyber risk management

the maritime equipment cybersecurity standard, since in both domains the artifact that is missing is not an analysis technique but a required place to write the answer down.

REFERENCES

[1] Cybersecurity and Infrastructure Security Agency, "CPDLC over ATN-B1 vulnerabilities," ICS Advisory ICSA-26-219-01, Aug. 7, 2026. [Online]. Available: https://www.cisa.gov/news-events/ics-advisories/icsa-26-219-01

[2] Cybersecurity and Infrastructure Security Agency, "FURUNO FA-50," ICS Advisory ICSA-26-237-07, Aug. 25, 2026. [Online]. Available: https://www.cisa.gov/news-events/ics-advisories/icsa-26-237-07

[3] M. Strohmeier, V. Lenders, and I. Martinovic, "On the security of the automatic dependent surveillance-broadcast protocol," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 2, pp. 1066–1087, 2015, doi: 10.1109/COMST.2014.2365951.

[4] Z. Wu, T. Shang, and A. Guo, "Security issues in automatic dependent surveillance-broadcast (ADS-B): A survey," *IEEE Access*, vol. 8, pp. 122147–122167, 2020, doi: 10.1109/ACCESS.2020.3007182.

[5] M. R. Manesh and N. Kaabouch, "Analysis of vulnerabilities, attacks, countermeasures and overall risk of the automatic dependent surveillance-broadcast (ADS-B) system," *Int. J. Crit. Infrastruct. Protect.*, vol. 19, pp. 16–31, Dec. 2017, doi: 10.1016/j.ijcip.2017.10.002.

[6] K. Sampigethaya and R. Poovendran, "Aviation cyber–physical systems: Foundations for future aircraft and air transport," *Proc. IEEE*, vol. 101, no. 8, pp. 1834–1855, Aug. 2013, doi: 10.1109/JPROC.2012.2235131.

[7] A. Gurtov, T. Polishchuk, and M. Wernberg, "Controller–pilot data link communication security," *Sensors*, vol. 18, no. 5, art. 1636, May 2018, doi: 10.3390/s18051636.

[8] J. Smailes, D. Moser, M. Smith, M. Strohmeier, V. Lenders, and I. Martinovic, "You talkin' to me? Exploring practical attacks on controller pilot data link communications," in *Proc. 7th ACM Cyber-Physical Syst. Security Workshop (CPSS)*, Virtual Event, Hong Kong, May 2021, pp. 53–64, doi: 10.1145/3457339.3457985.

[9] D. Alhalabi, M. Alawida, and R. Hasan, "How secure is CPDLC? A comprehensive review of cyber threats and defense mechanisms in aviation

communication," *J. Transp. Security*, vol. 18, art. 23, Oct. 2025, doi: 10.1007/s12198-025-00312-z.

[10] M. Ziazi, K. Aleem, H. Sathaye, and M. Strohmeier, "Sliding into the flight deck's DMs: Practical message attacks on CPDLC," in *Proc. 35th USENIX Security Symp.*, 2026. [Online]. Available: https://www.usenix.org/conference/usenixsecurity26/presentation/ziazi

[11] M. Balduzzi, A. Pasta, and K. Wilhoit, "A security evaluation of AIS automated identification system," in *Proc. 30th Annu. Computer Security Applications Conf. (ACSAC)*, New Orleans, LA, USA, Dec. 2014, pp. 436–445, doi: 10.1145/2664243.2664257.

[12] "AIS cybersecurity: Challenges, vulnerabilities, and mitigation strategies," *J. Marine Sci. Eng.*, vol. 14, no. 14, art. 1258, Jul. 2026, doi: 10.3390/jmse14141258.

[13] A. Goudossis and S. K. Katsikas, "Towards a secure automatic identification system (AIS)," *J. Marine Sci. Technol.*, vol. 24, no. 2, pp. 410–423, Jun. 2019, doi: 10.1007/s00773-018-0561-3.

[14] G. C. Kessler, "Protected AIS: A demonstration of capability scheme to provide authentication and message integrity," *TransNav*, vol. 14, no. 2, pp. 279–286, Jun. 2020, doi: 10.12716/1001.14.02.02.

[15] "A pragmatic approach to VDES authentication," *NAVIGATION: J. Inst. Navigation*, vol. 72, no. 1, art. navi.681, 2025. [Online]. Available: https://navi.ion.org/content/72/1/navi.681

[16] M. Strohmeier, V. Lenders, and I. Martinovic, "Detecting malicious ADS-B broadcasts using wide area multilateration," in *Proc. IEEE/AIAA 34th Digital Avionics Syst. Conf. (DASC)*, Prague, Czech Republic, Sep. 2015, doi: 10.1109/DASC.2015.7311579.

[17] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection," *Proc. IEEE*, vol. 104, no. 6, pp. 1258–1270, Jun. 2016, doi: 10.1109/JPROC.2016.2526658.

[18] C. V. Ribeiro, A. Paes, and D. de Oliveira, "AIS-based maritime anomaly traffic detection: A review," *Expert Syst. Appl.*, vol. 231, art. 120561, Nov. 2023, doi: 10.1016/j.eswa.2023.120561.

[19] "New machine learning approaches for intrusion detection in ADS-B," arXiv:2510.08333, Oct. 2025. [Online]. Available: https://arxiv.org/abs/2510.08333

[20] S. Kriaa, L. Pietre-Cambacedes, M. Bouissou, and Y. Halgand, "A survey of approaches combining safety and security for industrial control systems," *Rel. Eng. Syst. Saf.*, vol. 139, pp. 156–178, Jul. 2015, doi: 10.1016/j.ress.2015.02.008.

[21] E. Lisova, I. Šljivo, and A. Čaušević, "Safety and security co-analyses: A systematic literature review," *IEEE Syst. J.*, vol. 13, no. 3, pp. 2189–2200, Sep. 2019, doi: 10.1109/JSYST.2018.2881017.

[22] G. Kavallieratos, S. Katsikas, and V. Gkioulos, "Cybersecurity and safety co-engineering of cyberphysical systems — A comprehensive survey," *Future Internet*, vol. 12, no. 4, art. 65, Apr. 2020, doi: 10.3390/fi12040065.

[23] N. Johnson and T. Kelly, "An assurance framework for independent co-assurance of safety and security," arXiv:1903.01220, Mar. 2019. [Online]. Available: https://arxiv.org/abs/1903.01220

[24] N. Johnson and T. Kelly, "Independent co-assurance using the safety-security assurance framework (SSAF): A Bayesian belief network implementation for IEC 61508 and Common Criteria," arXiv:2010.07288, Oct. 2020. [Online]. Available: https://arxiv.org/abs/2010.07288

[25] RTCA, *DO-326A: Airworthiness Security Process Specification*, Washington, DC, USA, 2014; and EUROCAE, *ED-202A*, Saint-Denis, France, 2014.

[26] EUROCAE, *ED-120: Safety and Performance Requirements Standard for Air Traffic Data Link Services in Continental Airspace*, Saint-Denis, France; and RTCA, *DO-290*, Washington, DC, USA.

[27] International Electrotechnical Commission, *IEC 63154:2021 — Maritime Navigation and Radiocommunication Equipment and Systems: Cybersecurity — General Requirements, Methods of Testing and Required Test Results*, Geneva, Switzerland, 2021.