TECHNICAL EXPLORATION REPORT

From AIS traffic screening to acoustic association

What the validated CTBTO 2026 exploration supports, rules out, and recommends next

Prepared for Rui, Igor, and Juan

Evidence cutoff: 15 August 2026 — integration commit d6aa721

Validated integrated project evidence only

Executive summary

BOTTOM LINE **The project can now describe retained AIS traffic and path-level transmission loss credibly. It cannot yet show a stable aggregate acoustic response to the public-AIS shipping screen, a repeated direction match, ship-specific attribution, or out-of-sample prediction.**

The scientific question is whether low-frequency energy measured at IMS hydrophone arrays can be explained or predicted by a chain from AIS traffic through a published ship source model and underwater propagation. The exploration deliberately advanced one claim rung at a time. That discipline prevented polished propagation products, nearby vessels, or visually suggestive spectra from being mistaken for acoustic attribution. [P1; P3]

- **Traffic characterization** is supported only for the exact retained public product. The data reveal strong distance and class selection, anonymous month-scoped identities, hourly thinning, blank identifiers, and 16 truncated monthly sources. They do not reveal a defensible percentage of real ships missing. [R5; P2]
- **The 678-row Juan cohort** is now a durable TL-only result at 63 and 80 Hz. RAM, spherical spreading, TL_lite, and tl_fast differ substantially, but the sign of RAM versus matched spherical spreading varies by path. The quick laws are useful screens, not corrections or final propagation. [R1; P1]
- **The full-February six-band test** is null or ambiguous under the frozen public-AIS/JOMOPANS/tl_fast estimator: weak same-band associations, high-minus-low contrasts, and adjusted slopes all have uncertainty intervals crossing zero, while future-exposure placebos are often comparable. This does not show that ships made no contribution. [R6]
- **The Heaney-inspired fixed-scene diagnostic** did not rescue the aggregate model. The full screen produced little peak/floor prominence; a nearest-range baseline was stronger, but that is a held-out proximity hypothesis rather than attribution. [R7]
- **The frozen Juan-table measured-acoustic follow-up gives limited event-level energy consistency: 3/4 modeled high/control pairs are positive at both 63 and 80 Hz, but 0/4 high scenes have strict directional support. It remains rung-2 only, not attribution. [R11]**
- **Real-event direction lanes did not clear their prerequisite quality gates. DIMIJOHN A, MAERSK DETROIT, and SEA DOLPHIN are indeterminate; the August acoustics-first search found three acoustic candidates but zero exact ARA rows compatible in both time and direction. None advances to RAM. [R8; R9; R10]**
- **The correct stopping decision** is to hold attribution and prediction. A defensible next phase requires materially better independent traffic coverage and/or a prospectively fixed corpus of repeated, complete-triplet real passages. More post-hoc tuning or a full 678-row environmental rerun would add cost without raising the claim ceiling. [P1; P3]

Current claim ceiling

Rung	Claim	Status	Evidence boundary
1	AIS traffic characterization	QUALIFIED	Supported for retained bytes and screenability; not actual traffic completeness.
2	High- vs low-exposure acoustic difference	NOT SUPPORTED	Full-February six-band result is null/ambiguous under the tested screen. The later four-pair Juan event follow-up is energy-only rung-2 consistency, with no directional support or claim advance.
3	Repeated AIS-bearing / triplet direction consistency	NOT SUPPORTED	Exact events are indeterminate; August search has zero survivors.
4	Ship-specific attribution	STOP	Prerequisite rungs are not met.
5	Out-of-sample prediction	STOP	No claim-level basis for model fitting or validation.

Evidence snapshot

Scale	Propagation / aggregate	Scale	Sensitivity / events
678	ordered H01W1 RAM paths at both 63 and 80 Hz	526	late-month hours in the primary six-band catalogue
12	fixed paths in the ocean temporal follow-up	3 -> 0	August acoustic candidates -> exact-row ARA survivors

The report reads from the strongest completed products toward the association bottleneck: propagation first, then environment and traffic representation, then measured aggregate and exact-event evidence. Nulls and indeterminate results remain visible because they determine the stopping decisions.

Decision map - what moves, what holds, what stops

This page translates the evidence into operating decisions. A lane moves only when its stated prerequisite is met; no negative or indeterminate result is converted into evidence of vessel absence.

Lane	Decision	Use now	Hard boundary
Public AIS	KEEP	Use for scene-level traffic description and broad screening.	Never report one completeness or missing-ships percentage.
Source model	KEEP	Use published JOMOPANS-ECHO as the visible nominal reference.	Keep its scatter and below-20-Hz uncertainty explicit.
Propagation	KEEP	Use tl_fast only to rank; use exact 2-D RAM for promoted paths.	No fitted quick-law offset; no dense fan or 3-D campaign now.
Ocean state	HOLD	P1M for broad screen; P1D for promoted-event sensitivity; PT6H secondary.	No global correction and no full 678-row rerun.
Aggregate acoustics	HOLD	Preserve the six-band null/ambiguous result.	Reopen only with materially better independent traffic/background data or a distinct frozen model.
Exact events	STOP	Close DIMIJOHN A, MAERSK DETROIT, SEA DOLPHIN, and the three August candidates as tested.	New events need complete triplets, repeated passages, and Lane A quality before Lane B/RAM.
Attribution/ prediction	STOP	Do not claim or fit.	Requires rung 2 and repeated rung 3 evidence first.

Priority next steps

1. Secure an independent, higher-cadence or otherwise better-characterized traffic source for a frozen H01 replication period, with known coverage and identity semantics.
2. Predefine a repeated-passage corpus with complete observed triplets; run exact-spherical Lane A first and stop every event that fails pair confidence and closure.
3. For the few promoted paths only, retain exact RAM and add P1D/PT6H as sensitivity representations; do not choose the representation by acoustic agreement.
4. Treat background and biologics as explicit competing explanations. February is lower concern than the main local blue/fin-whale seasons, but it is not whale-free and 63/80 Hz energy alone is non-diagnostic.

1. Scientific objective and evidentiary architecture

The target is not a transmission-loss atlas. It is evidence for a chain: AIS type, dimensions, speed, position, and time; a published source spectrum; path transmission loss; modeled individual or aggregate ship contribution; and a measured hydrophone response, ideally with compatible direction on a complete triplet. Each link can fail for a different reason, and a strong result in one link does not substitute for the next. [P1; P3]

A vocabulary that prevents false equivalence

Quantity	What it means here
Source spectral or band level (SL)	A modeled ship emission reference, usually at 1 m and for a stated band. It is not measured at the hydrophone.
Transmission loss (TL)	Propagation loss along a specified path, frequency, environment, and depth convention. It contains no ship source amplitude.
Individual received level (RL_i)	A modeled contribution from one ship: SL - TL, under the same reference conventions.
Modeled shipping-only aggregate (NL_ships)	The linear-power sum of unrelated modeled RL_i values. It excludes unrepresented ships and non-shipping background.
Measured ambient energy	Hydrophone energy containing shipping, biologics, weather, geophysics, instrumentation, and other sources. Some products are nominally response-corrected; others are relative raw-count power.
Screening prediction	A rough AIS/JOMOPANS/range-law quantity used to rank scenes. It is not a calibrated prediction of measured ambient level.
Attribution	A repeated, quality-controlled inference that a particular vessel contributed to a measured feature, supported by timing, direction, spectral/context evidence, and alternatives.

INTERPRETIVE RULE **A modeled RL_i below a measured ambient level is not automatically a failure: measured ambient energy includes background and other sources. Conversely, a nearby AIS row plus an acoustic peak is not attribution.**

2. Data landscape and its limiting asymmetries

The evidence base combines broad but anonymized public traffic data, sparse identified ARA snapshots, restricted local acoustic holdings, and environmental fields with different temporal meanings. These sources are complementary rather than interchangeable.

Evidence layer	Scope	Strength	Limitation
AMSA public monthly traffic	19 HA01 extracts; only 3 source months structurally complete	Broad scene coverage; type/length/speed for many fixes	At least hourly thinning; anonymous month-scoped CRAFT_ID; 16 readable prefixes; no named-vessel proof
Juan selected cohort	678 rows in 233 UTC hour groups; 674 source-model eligible	A reproducible selected screening cohort across Jan 2025-Jul 2026	Selection-heavy; 578 rows from truncated source months; geometry originally absent
ARA identified snapshots	Irregular overlapping latest-known-position censuses	Exact identity and point-row event hypotheses	Not a continuous feed; row ages and provider/cadence/field semantics incomplete
Restricted H01 acoustics	Full February triplet hourly products; August package with extensive missingness	Measured same-band energy and bounded triplet direction tests	Calibration/response limits; channel health is event-specific; restricted data remain local
GEBCO + Copernicus ocean	Path bathymetry plus P1M, P1D, and instantaneous PT6H fields	Reproducible RAM environments and sensitivity tests	Model representations, not truth; limited coincident observations

A station being publicly listed as certified does not prove that every hydrophone channel was healthy in a target window. The H01 and H10 exact-event lanes therefore require observed triplet status and event-specific quality checks. Likewise, public AIS absence is not evidence of vessel absence. [P2; P3]

3. Juan's 678-row propagation result

Juan's handoff became a complete, ordered H01W1 TL table at nominal 63 and 80 Hz. All 678 rows retain their workbook identity and exact WGS84 geometry. The 63 Hz RAM values span 72.280-127.260 dB with median 100.125 dB; 80 Hz spans 70.016-131.042 dB with median 98.267 dB. RAM80 - RAM63 is mixed: mean -2.007 dB, median -1.754 dB, and p10/p90 -9.953/+6.078 dB. Frequency therefore cannot be represented by a fixed offset. [R1; P2]

The comparison is deliberately TL-only. It does not import a source level, compute received level, open a measured acoustic outcome, tune a model, or rank contribution. This boundary matters because Juan's historical columns named NL and NLsum are screening quantities derived from a selected cohort and a range law; they are not measured ambient noise. [R1; P1]

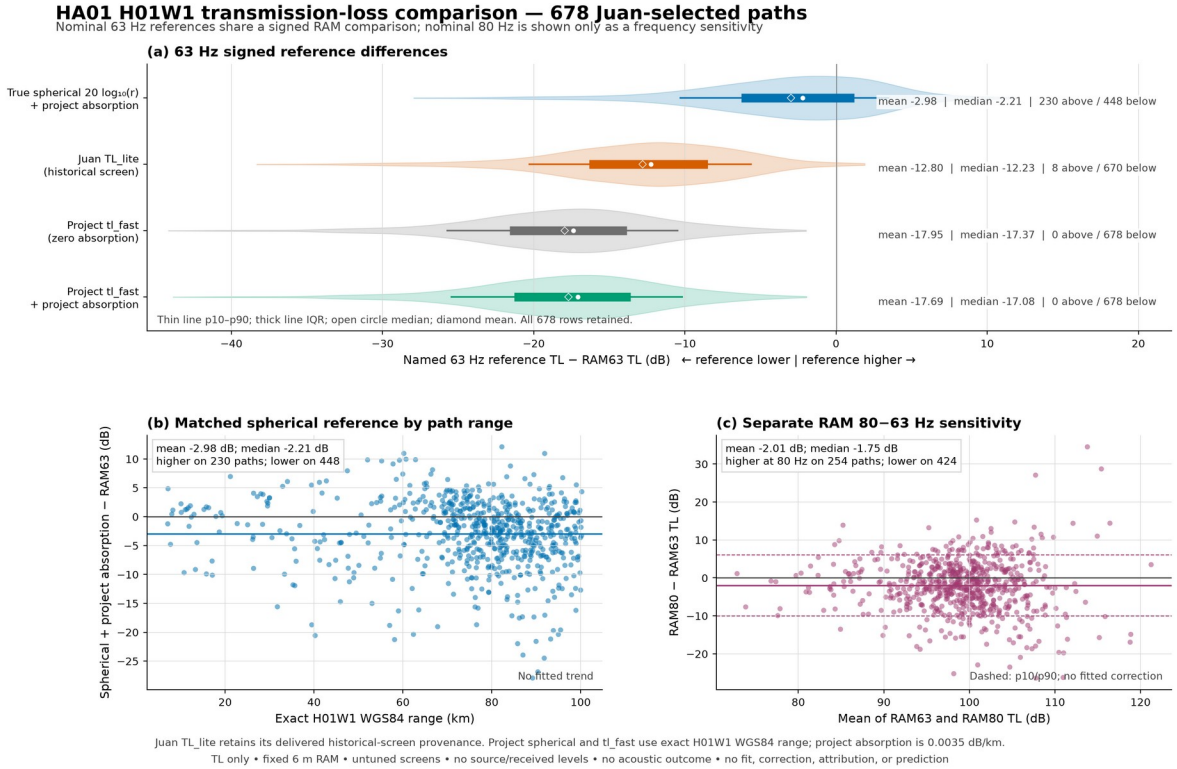


Figure 1. Full-cohort transmission-loss comparison for all 678 Juan-selected paths. The mixed spreading laws sit below RAM for nearly all paths, whereas matched spherical spreading crosses zero; the 80-63 Hz RAM difference is also two-sided.

Provenance: [R1] data/derived/ha01-juan-tl-comparison/v1/ha01-juan-tl-comparison.png

Quick-law bias: what it does and does not mean

Across the full cohort, reference-minus-RAM63 means are -2.978 dB for matched spherical spreading plus project absorption, -12.804 dB for Juan TL_lite, and -17.689 dB for project-absorption tl_fast. Spherical is higher than RAM on 230 paths and lower on 448; the quick laws are lower on almost every path. [R1]

The correct interpretation is not that RAM is always more lossy than spherical spreading. The quick laws contain long sub-spherical 10-log and 15-log branches, which create a structural downward regime bias

relative to a coherent solver. In the separately frozen 18-path benchmark, those branches account arithmetically for most of the mean gap, while matched-absorption RAM versus true spherical has no stable sign. [P1; P3]

A leading pathwise mechanism is coherent sampling at the fixed 6 m reference depth, roughly one quarter wavelength below a pressure-release surface at 63 Hz. Surface-image cancellation and weak near-surface modal amplitudes can move individual RAM values to either side of spherical spreading. This explanation is physically plausible and qualitatively supported by full-depth fields, but there is no clean path-level decomposition and no basis for a fixed 6 m correction. [P1; P3]

DECISION **Use tl_fast as an untuned ranking and sanity-check model. Retain TL_lite for historical provenance. Use exact RAM for promoted paths, and fit no offset from these comparisons.**

4. Source depth is a convention before it is a sensitivity

Production remains at a fixed 6 m nominal/equivalent-monopole reference because the adopted JOMOPANS-ECHO source levels are standardized to that depth. In reciprocal RAM runs the numerical source remains at the hydrophone depth and the ship-side field is sampled at 6 m. This is an alignment of reference conventions, not a claim that machinery, propeller, cavitation, or distributed ship sound is physically concentrated at 6 m. [R2; S2]

The recovered 3-12 m propagation-only sweep held source level fixed and reproduced both accepted 6 m paths exactly. Maximum conditional RL_i change from 6 m was 5.734 dB for MAERSK DETROIT and 5.810 dB for CRETANSEA - moderate under the frozen categories and close to, but below, the approximate 6 dB JOMOPANS source-model scatter. Because source level was not re-referenced, these are sensitivity curves, not complete alternative source conventions. [R2]

The later authorized extension to the frozen 18-path benchmark is a permanent non-result. Its final attempt stopped after one RAM call retained an opaque, non-accepted raw field and before any accepted TL extraction; zero rays were accepted, 15 were not run, and no table, plot, correction, or retry follows. The validated two-path H10 sweep remains the only numerical depth-sensitivity evidence. [R12]

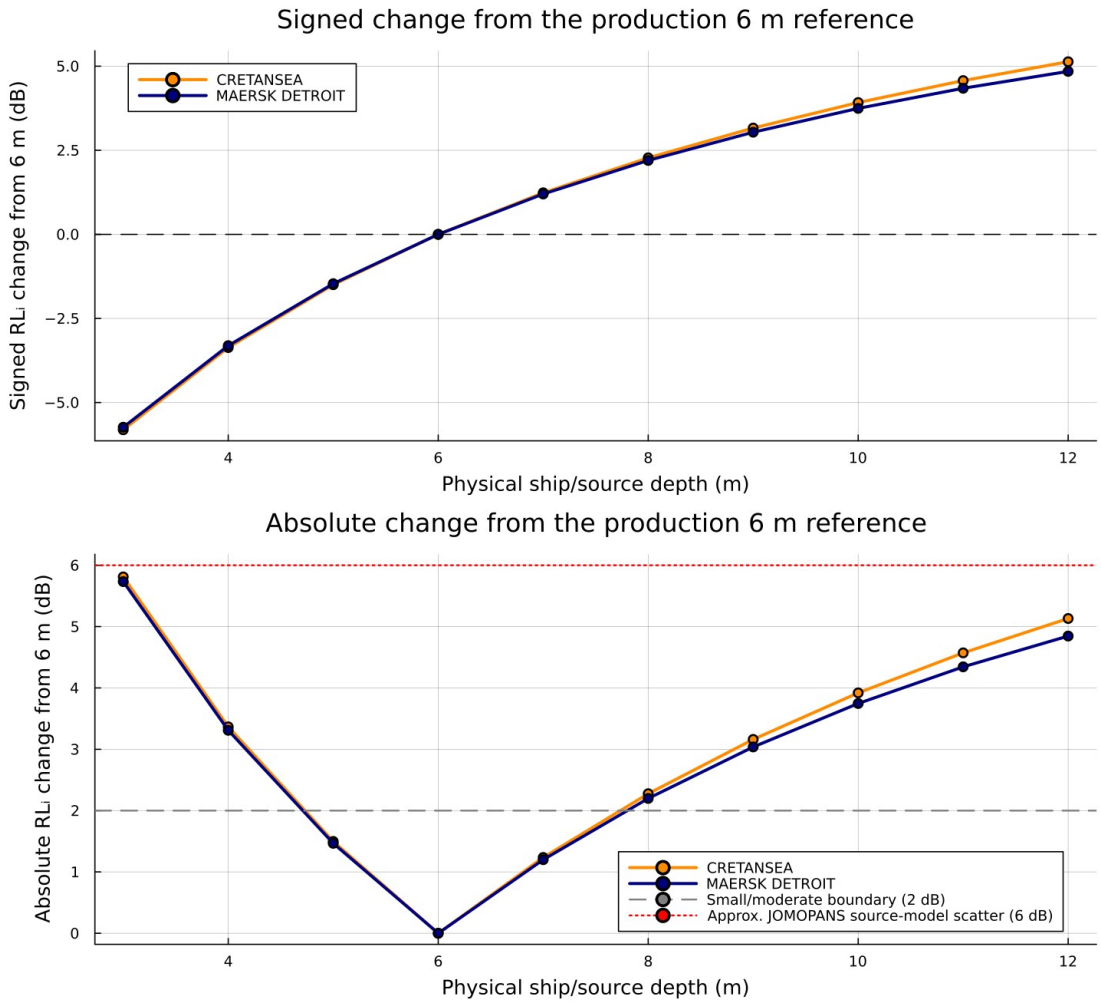


Figure 2. Conditional change from the fixed 6 m production reference for two accepted H10 paths. Depth sensitivity is material but remains below the predeclared 6 dB large-effect threshold.

Provenance: [R2] data/derived/h10-source-depth-sensitivity-3m-12m-v1/change-from-6m.png

A genuine comparison among 0.5 x draught, 6 m, and 0.7 x draught would need matched re-referencing of both source level and propagation depth, fixed before any acoustic comparison. Repeating propagation-only depth sweeps does not resolve that convention question. [R2; P3]

5. Exact environment handoff: GEBCO and sound-speed profiles

The exact as-run environment for all 678 paths was exported without rerunning propagation: 51,027 bathymetry knots, 5,717 range-dependent SSP profiles, 320,630 finite sound-speed values, and 35 deduplicated depth axes. The ordered workbook companion preserves all original 63/80 TL cells and adds two portable references: simple GEBCO water depth at the exact AIS fix and the ship-end SSP filename. [R3]

The simple AIS-fix GEBCO depth is a reference, not the bathymetry used by RAM. RAM used the entire range-dependent path and sound-speed field. Requested sampling near 1 km for bathymetry and 10 km

for SSP determined counts; actual endpoint-inclusive spacing depends on path length. Reports should therefore preserve actual knot ranges rather than imply exact uniform 1 km/10 km grids. [R3; P1]

PRACTICAL VALUE **The export makes Juan's 63/80 TL table inspectable and portable while preserving the difference between a single depth-at-ship lookup and the full environment that generated RAM TL.**

6. Ocean temporal resolution changes TL path by path

The completed 12-path H01 follow-up compared three temporal representations of the same Copernicus product at nominal 63 Hz, holding geometry, bathymetry, bottom proxy, numerics, 6 m depth, and extraction fixed. Every monthly solve reproduced the accepted Juan-production value exactly. [R4]

Contrast	Threshold counts	Median |change|	Range	Sign balance
Daily - monthly	9/12 at >=1 dB; 4/12 at >=3 dB	1.917 dB	-4.955 to +8.459 dB	6 positive / 6 negative
PT6H - daily	2/12 at >=1 dB; 1/12 at >=3 dB	0.220 dB	-3.368 to +1.050 dB	6 positive / 6 negative

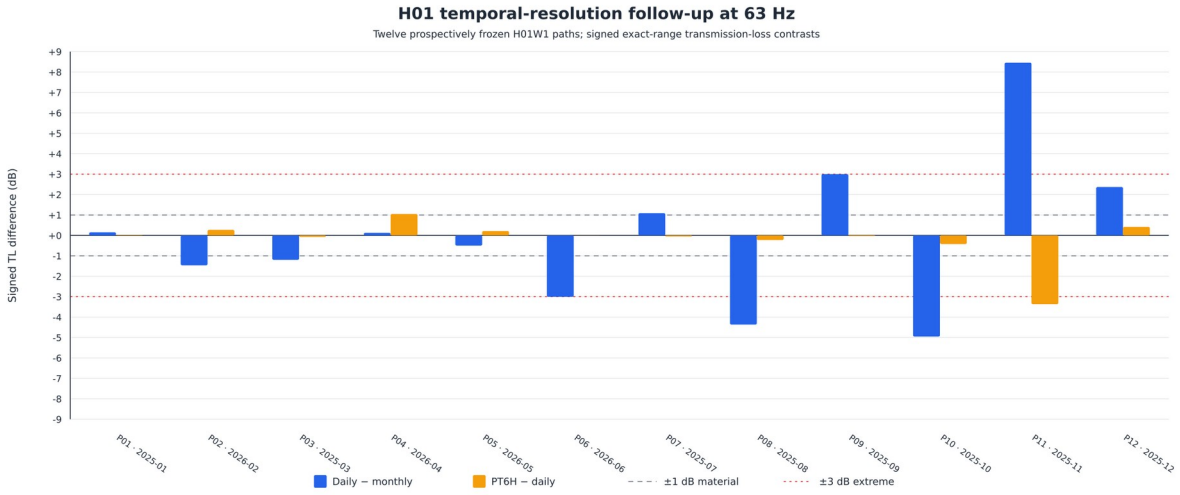


Figure 3. Signed TL sensitivity to monthly, daily, and nearest instantaneous six-hourly ocean fields. The current 12-path follow-up has mixed signs and isolated large changes, so no global correction is defensible.

Provenance: [R4] data/derived/ha01-ocean-temporal-followup/v1/temporal-deltas-63hz.svg (rasterized for Word)

Physical and observational interpretation

P1M is a UTC calendar-month mean and remains an economical broad-screen baseline. P1D is an exact UTC-day mean and is preferred for promoted event sensitivity because it is more temporally specific, not because it has been proven correct. PT6H is an instantaneous field at six-hour cadence, not a six-hour mean and not an estimate of a diurnal cycle. [R4; R5]

Public observations did not provide a qualifying subsurface coincidence for the two predeclared extreme paths. The best inspected Argo profile for one case was about 46.7 km away and 41.3 hours later; the other had no useful nearby profile. Regional upper-ocean variability is plausible, but it cannot explain the sign or magnitude of a coherent endpoint change without coincident full-depth observations and a field-fragility diagnostic. [R5]

DECISION **Keep P1M for the broad 678-row screen. Use P1D, with PT6H as a secondary sensitivity, only for promoted events. Fit no correction and do not rerun all 678 paths.**

7. AIS representation is measurable; missing ships are not

Across the 19 retained AMSA months, 16 sources are readable prefixes and only three are structurally complete. The known national structural shortfall is 2,941,988 of 49,897,448 declared records (5.90%), but the spatial distribution of the missing tails is unknowable. The 678-row selected cohort represents only 0.42% of 159,925 visible HA01 rows, includes 578 rows from truncated months, and contains no selected rows from the visible 150-200 km ring. [P1; P2]

The complete-month waterfall answers a narrower, defensible question: what do the exact retained public bytes represent, and what fraction can enter successively narrower screens? July 2026, June 2026, and February 2025 contain 5,712, 6,198, and 12,221 fixes within 200 km of the H01W centroid; exact JOMOPANS eligibility is 76.2%, 71.6%, and 54.0%. [R5]

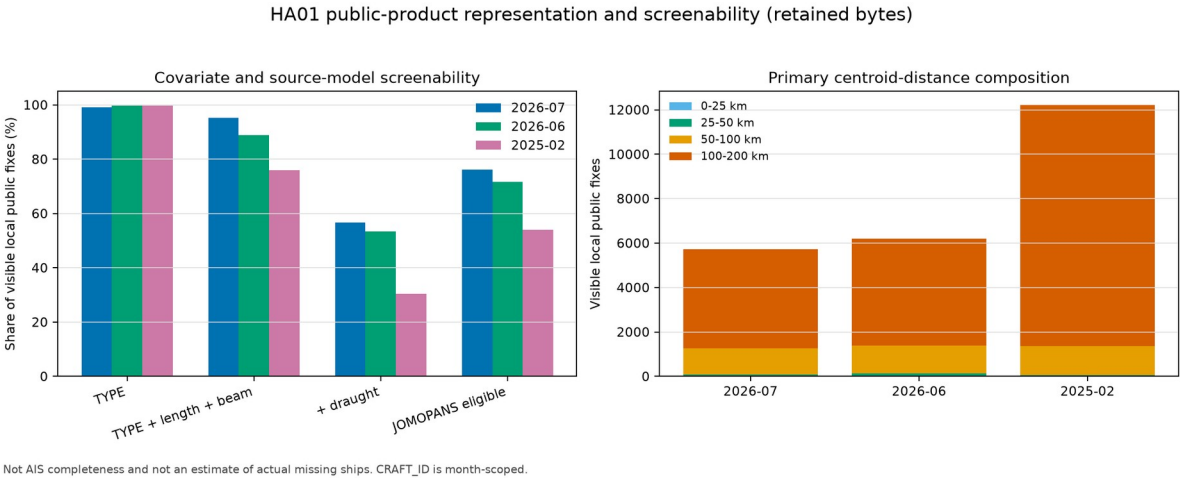


Figure 4. Public-product representation and source-model screenability in the three structurally complete months. Most retained fixes lie 100-200 km from H01; covariate availability narrows the usable population.

Provenance: [R5] data/derived/ha01-ais-representation-waterfall/v1/representation-waterfall.png

June illustrates why one completeness percentage would be false: 4,647 fixes lack public CRAFT_ID. Those rows still count as visible fixes and may be source-model screenable, but they cannot support same-ID trajectories, craft-hours, passages, or empirical cadence. Even for nonblank identifiers, CRAFT_ID is anonymous and month-scoped, not MMSI or IMO. [R5]

The first route-interpolation calibration also failed its all-part gate: a route-conditioned prior improved the p90 position error but worsened the median and every outer-month median. The supported coarse-position baseline remains great-circle interpolation with held-out empirical uncertainty when two-sided within-month fixes exist. It does not reconstruct a ship track. [P1; P3]

8. Full-February aggregate acoustic test

The six-band analysis froze exposure before opening any acoustic value. It used 23,192 retained public AMSA fixes within 300 km, JOMOPANS-ECHO source levels, exact routes to the three H01W nodes, and untuned tl_fast screens. The measured response was full-hour Welch power in nominally response-corrected Pa^2, averaged in linear power across the triplet before conversion to dB. [R6]

The least-outcome-exposed late-month catalogue contains 526 hours. Across 10, 16, 25, 40, 63, and 100 Hz, same-band Spearman estimates are weakly positive (0.016-0.120), but every 2,000-day-cluster bootstrap interval crosses zero. High-minus-low exposure-quartile contrasts range from -0.184 to +0.312 dB and every interval crosses zero; adjusted slopes are near zero. Receive-time and uniform-attenuation variants do not change the result. [R6]

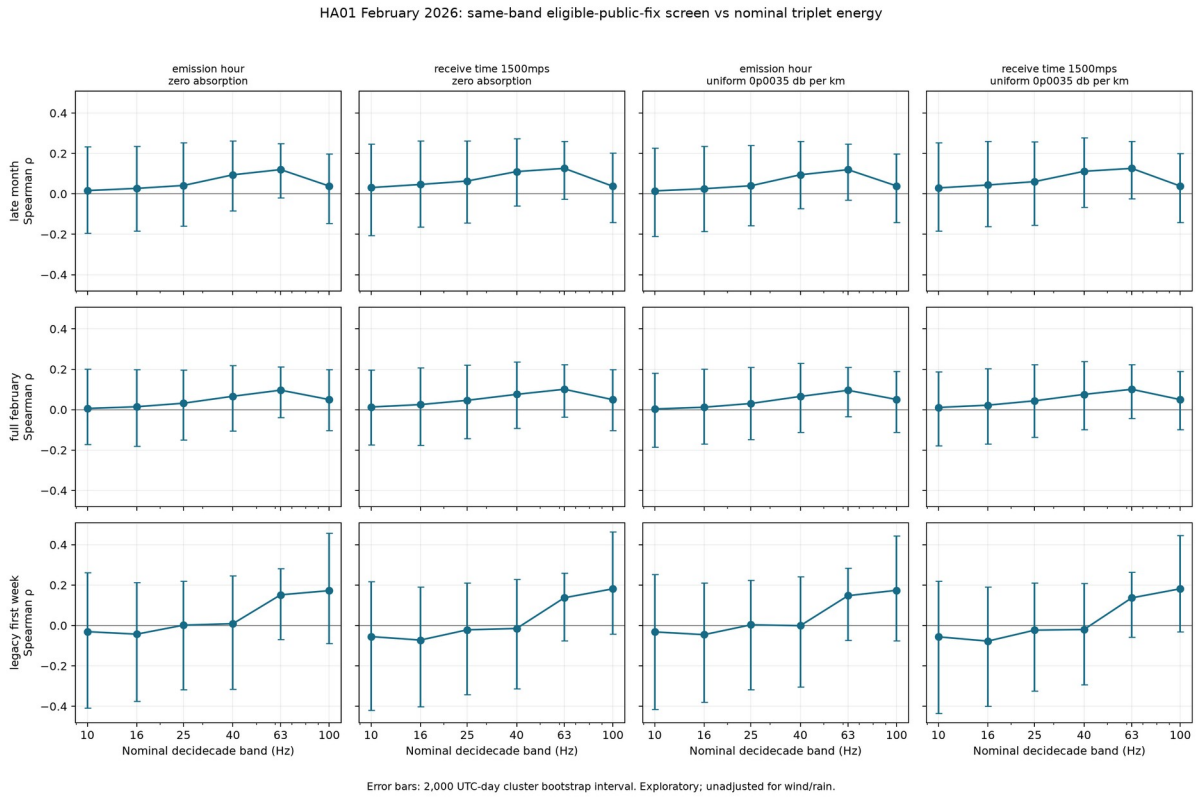


Figure 5. Same-band modeled exposure versus measured triplet energy across the frozen February views. Point estimates are small, uncertainty intervals cross zero, and the pattern is not robust to placebo comparison.

Provenance: [R6] data/derived/ha01-feb-six-band-aggregate/v1/six-band-association-summary.png

Future-exposure placebos at +6 and +24 hours are often comparable to or stronger than the physically aligned estimates. The screen also does not consistently outperform simple source-level, craft-count, or nearest-range baselines. This is no stable aggregate shipping-energy association under the tested

estimator. It is not evidence against a ship contribution: public-AIS thinning and truncation, unrepresented ships, propagation error, masking, biologics, and missing wind/rain adjustment remain live explanations. [R6]

Bounded Juan-table measured-acoustic follow-up

A separate frozen pass over the existing monthly-P1M RAM63/RAM80 table found 40 of 678 rows in 12 February 2026 hours with validated complete H01W triplet coverage. Four non-overlapping same-month anonymous public-AMSA-ID high/control pairs were selected before acoustic outcomes. Controls have lower modeled contribution; they are not quiet, environmentally matched, or named identities. [R11]

Three of the four pairs had positive measured high-minus-control contrasts at both nominal 63 and 80 Hz. No high scene passed the strict exact-node pooled 10-80 Hz direction rule; pair 3 is directionally disfavored for that fixed estimator and the other three are indeterminate. This is energy-only event-level rung-2 consistency, not detection, identity, attribution, causal evidence, RAM validation, calibrated absolute agreement, or prediction. [R11]

9. Heaney-inspired fixed-scene diagnostic

Heaney et al. supplied a useful forward-model diagnostic idea: freeze modeled scene times, then compare modeled and measured structure without tuning to agreement. The project kept the important limitation visible: the reference study tuned sediment using the same interval later plotted, so it is a feasibility template rather than an independent attribution design. [S1; P4]

Eight model-only scene catalogues were frozen before opening existing February hourly products. The four full public-AIS/JOMOPANS/tl_fast catalogues selected the same 12 scenes and produced median local peak/floor prominence of +0.055 dB, with a 95% scene-bootstrap interval from -0.296 to +0.532 dB and a 76.8th-percentile time-shift-null position. Only three of six band medians were positive. [R7]

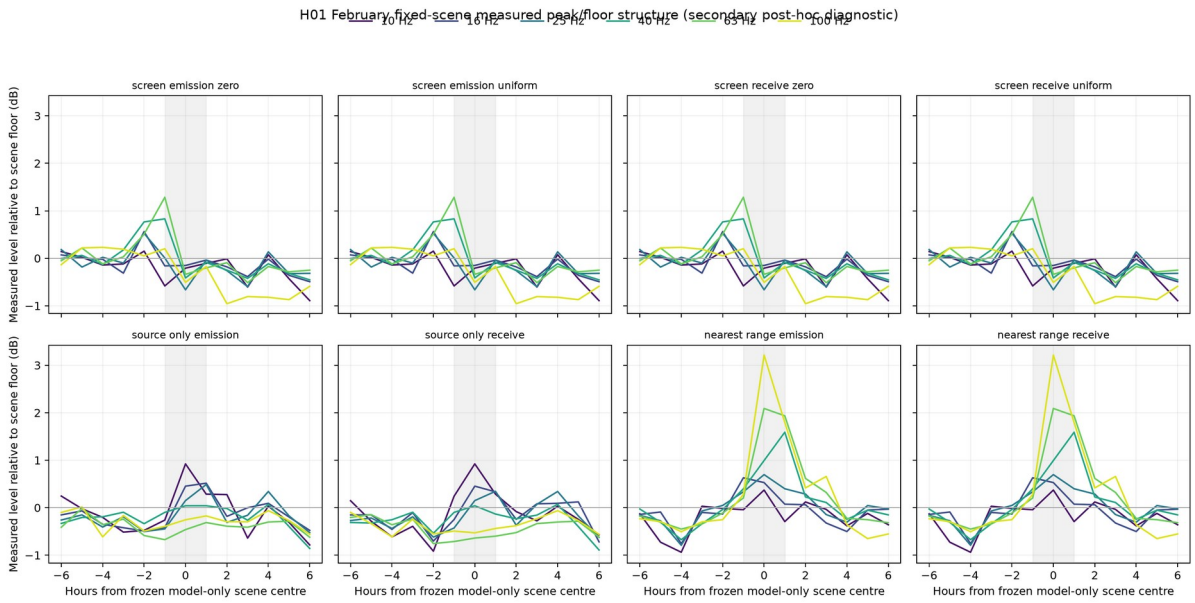


Figure 6. Measured peak/floor structure around eight fixed model-only scene catalogues. The full screen shows little central prominence; the nearest-range baseline is visibly stronger but remains a separate proximity hypothesis.

Provenance: [R7] data/derived/h01-heaney-fixed-scene/v1/fixed-scene-peak-floor.png

The nearest-range baseline was stronger: median prominence +0.989 dB, bootstrap interval +0.318 to +1.828 dB, and 100th percentile of its fixed shift null. That finding is worth preserving as a held-out proximity hypothesis, but it does not validate the full aggregate model, establish direction, identify a vessel, or overturn the six-band null/ambiguous result. [R7]

STOPPING RULE **Do not tune sediment, source depth, source level, propagation, scene centres, or bands to improve the overlay. Reopen only as a separately frozen held-out replication with better independent traffic coverage.**

10. Exact-event direction lanes

Exact-event tests ask a different question from hourly association: does a complete triplet contain a quality-valid coherent signal whose relative delays are compatible with an AIS hypothesis? Pair confidence and closure are prerequisites. A visually nearby scan maximum is non-evidence when those gates fail. [P3]

Lane	Frozen scope	Decisive result	Outcome	Decision
DIMIJOHN A, H10N	Exact event + fixed +24 h comparator	Both complete triplets; event 0/3 confidence passes, closure 0.553 s; comparator 0/3, closure 0.156 s	INDETERMINATE	Lane B not run; interval closed
MAERSK DETROIT, H10N	Exact event + fixed +24 h comparator	Both complete observed triplets; event 0/3 confidence, 1.562 s closure; comparator 0/3, 0.948 s	INDETERMINATE	Lane B not run; intervals closed
SEA DOLPHIN,	Two exact fixes + fixed	Both fail direction quality; 7/11 energy	INDETERMINATE	No SL/RL or RAM; lane

Lane	Frozen scope	Decisive result	Outcome	Decision
H01W	-24 h comparators	bands positive at each fix, but controls unmatched		closed
August acoustics-first, H01W	8,433 eligible 20.48 s windows	3 acoustic candidates; 691 unique exact ARA rows; zero rows compatible in time and direction	ZERO SURVIVORS	No RAM or event escalation

DIMIJOHN A

The complete H10N triplet was structurally valid at the exact event and +24-hour comparator, but neither window had a reliable delay solution. Poor agreement with AIS delay, scan bearing, or null ranking therefore cannot be interpreted against the vessel. The correct result is indeterminate, not detected, absent, contributing, or non-contributing. [R8]

MAERSK DETROIT

A second outcome-blind H10N exact-event test selected MAERSK DETROIT first among 19 eligible cargo rows. The 20.48-second event and fixed +24-hour comparator were complete observed triplets, but both failed the Lane A prerequisite: event 0/3 confidence passes with 1.562262 s closure against the 0.004 s limit, comparator 0/3 with 0.948177 s closure. The outcome is indeterminate, not directional disfavor. H10S was excluded because no complete target-window triplet was available. No Lane B, energy/amplitude, calibration, RAM, propagation, attribution, or prediction work ran; both intervals are closed. [R13]

SEA DOLPHIN

The close fix passed only two of three pair-confidence checks and exceeded the 4 ms closure limit; the later fix passed none and had a very large closure error. Event-minus-comparator energy was positive in seven of eleven bands at each fix, especially from 31.5-100 Hz, but the comparators were not verified low-exposure or environmentally matched. Without quality-valid direction, this descriptive amplitude pattern cannot support a ship contribution claim. [R9]

Completed August acoustics-first search

The acoustics-first membership was frozen before opening ARA snapshots. Three candidates survived fixed coherence, energy, direction-stability, separation, and ranking gates. After reveal, one candidate had seven exact ARA rows within +/-30 minutes but the best directional residual was 132.669 degrees; the other two had no exact rows in the time window. There were zero rows jointly compatible in time and direction, zero repeated compatible identities, and zero survivors. [R10]

This is negative compatibility evidence against the sparse exact-row ARA snapshots, not evidence that the acoustic candidates were not ships or that no vessel contributed. The snapshots are irregular censuses rather than a complete continuous track source. No interpolation, closest point of approach, travel-time adjustment, source model, or RAM was invoked. [R10]

11. Whale-season caveat

Whale activity affects interpretation of measured acoustics, not TL. Cape Leeuwin literature places Southeast Indian Ocean pygmy blue-whale detections mainly in March-May with a smaller December peak, Antarctic blue whales mainly in May-August, and fin whales mainly in April-October. February is therefore lower concern than May-June for 63/80 Hz masking, but it is not guaranteed whale-free. No direct February 2026 biological detector was run. Blue-whale harmonics can extend into these bands; band elevation alone cannot establish whale presence. [S6-S8; P4]

USE **Carry biologics as a competing-background flag in measured-energy analyses. Do not use seasonality to erase an acoustic feature, and do not infer whale presence from 63/80 Hz energy alone.**

12. What is supported, not supported, and learned

Category	Current synthesis
Supported	Retained-public-product traffic description; exact row geometry; 678-row 63/80 Hz TL; reproducible GEBCO/SSP environments; path-specific depth and ocean sensitivities; bounded candidate searches.
Not supported	A stable February high-versus-low acoustic effect; repeated AIS-direction consistency; ship-specific attribution; proof of vessel absence; global TL correction; calibrated absolute model-to-ambient agreement; out-of-sample prediction.
Learned	Quick laws have structural regime bias but useful ranking value; coherent near-surface sampling and ocean time representation create mixed pathwise variation; public AIS is selection- and representation-limited; quality gates prevent attractive but invalid event stories.
Still uncertain	True local missing traffic; ARA coverage/cadence semantics; event-specific background; source directivity/loading/cavitation; full-depth ocean truth; real-passage performance of direction and extraction methods.

The exploration has therefore been scientifically productive even without a positive attribution result. It converted several tempting shortcuts into explicit boundaries: a TL map is not an acoustic finding; NL labels in a screening table are not measured ambient noise; a nearby named ship is not a detection; a good-looking bearing is not direction evidence if closure fails; and absent sparse AIS rows do not prove vessel absence.

The most important positive result is methodological coherence. The project now has a defensible funnel from broad traffic characterization through source and propagation screening, exact RAM for finalists, measured energy, triplet direction, and only then attribution. The current evidence stops in that funnel exactly where it should.

13. Recommendations and explicit stopping decisions

1. Preserve the existing 678-row 63/80 TL tables and environment handoff as final for their stated scope. Do not reinterpret them as source level, received level, measured acoustics, or attribution.
2. Retain the fixed 6 m JOMOPANS-aligned convention. Treat 3-12 m as propagation uncertainty; compare 0.5T/6 m/0.7T only through a prospectively matched source-level and propagation-reference study.
3. Use tl_fast only for broad ranking and dominance screening. Do not fit an offset to RAM; run exact 2-D RAM only for candidates that first survive independent acoustic and traffic gates.
4. Keep P1M as the production screening environment. For promoted paths, use P1D and nearest PT6H as fixed sensitivities. Do not infer a global correction or rerun the whole cohort.
5. Treat the February six-band and Heaney-style results as closed null/ambiguous analyses under the tested estimator. Reopen only with materially better independent traffic/background coverage or a prospectively distinct model.
6. Do not reopen the exact DIMIJOHN A, MAERSK DETROIT, or SEA DOLPHIN intervals, and do not propagate the three August acoustics-first candidates. New event work should use a predeclared repeated-passage corpus with complete observed triplets.
7. Keep Lane A quality first. Only a Lane A pass permits fixed Lane B extraction; only repeated compatible events with defensible timing, background, and propagation can approach ship-specific attribution.
8. Do not begin out-of-sample prediction. Rung 2 and repeated rung 3 evidence are prerequisites, not optional covariates.
9. Report nulls and indeterminate results as results. They constrain methods and data requirements without implying that ships did not contribute.
10. Do not acquire high-cadence identified AIS yet. First obtain separate authorization to audit ARA provenance/licence and request only a 24-48 hour non-study sample. If that route passes, use the fixed 1-8 July 2026 historical option or, preferably, a prospectively synchronized seven-day campaign. Freeze six passages across at least three identified vessels and two dates plus six matched controls; require at least 4/6 two-band energy-consistent pairs and directional support in at least two different vessels/dates. Even a full pass is capped at rung 3, not attribution or prediction. [R14]

RECOMMENDED PROGRAM Do not begin a new AIS acquisition now. The next decision, if separately authorized, is an ARA provenance/licence audit plus a 24-48 hour non-study sample; only a passing preflight can release the identified-AIS replication design. One better-observed replication period and one prospectively fixed repeated-passage corpus are more valuable than a larger post-hoc search, a denser propagation fan, or a tuned correction.

Evidence and provenance map

Bracketed codes in the report point to the current integrated project artifacts below. Paths are relative to /home/rui/ara/ctbto-2026 unless stated otherwise. Body text omits implementation and runtime history; the referenced manifests retain those details.

[P1] Project state - 2026-08-15. Integrated current conclusions, decisions, and claim boundaries. Local: docs/PROJECT_STATE.md

[P2] Data inventory. Dataset scope, cardinalities, provenance, and limitations. Local: docs/DATA_INVENTORY.md

[P3] Study design. Claim ladder, estimands, gates, and interpretation rules. Local: docs/STUDY_DESIGN.md

[P4] Exploration portfolio - 2026-08-14. Heaney, source depth, environment, AIS, and whale-season synthesis. Local: docs/EXPLORATION_PORTFOLIO_2026-08-14.md

[R1] HA01 Juan-selected TL comparison. All 678 RAM/spherical/TL_lite/tl_fast rows and 63/80 sensitivity. Local: data/derived/ha01-juan-tl-comparison/v1/

[R2] H10 source-depth sensitivity and decision. Validated 3-12 m sensitivity and 6 m convention. Local: data/derived/h10-source-depth-sensitivity-3m-12m-v1/; docs/H10_SOURCE_DEPTH_DECISION.md

[R3] Juan RAM environment export. Exact as-run GEBCO/SSP handoff without propagation rerun. Local: data/derived/ha01-juan-ram-environment-export/v1/; analysis/juan-environment-workbook/v1/

[R4] H01 ocean temporal follow-up. 12-path monthly/daily/PT6H TL sensitivity. Local: data/derived/ha01-ocean-temporal-followup/v1/

[R5] H01 ocean observational validation and AIS representation. Product semantics, observation limits, retained AIS screenability. Local: analysis/ha01-ocean-observational-validation/v1/; data/derived/ha01-ais-representation-waterfall/v1/

[R6] HA01 February six-band aggregate. Frozen same-band aggregate association and placebos. Local: data/derived/ha01-feb-six-band-aggregate/v1/

[R7] H01 Heaney fixed-scene diagnostic. Eight frozen scene catalogues and peak/floor result. Local: data/derived/h01-heaney-fixed-scene/v1/

[R8] H10N DIMIJOHN A exact-event direction. Quality-gated indeterminate Lane A result. Local: analysis/h10n-dimijohn-direction/v1/

[R9] H01W SEA DOLPHIN exact-event lane. Two-fix direction and same-band contrast result. Local: analysis/h01w-sea-dolphin-exact-event/v1/

[R10] H01W August acoustics-first search. Integrated three-candidate, zero-survivor search result. Local: analysis/h01w-august-acoustics-first/v1/results.md

[R11] Juan measured-acoustic follow-up. Frozen four-pair 63/80 energy and exact-node direction result. Local: analysis/ha01-juan-measured-acoustic-followup/v1/

[R12] TL-screening benchmark source-depth extension. Permanent stop with no accepted depth result. Local: analysis/tl-screening-ram-benchmark/v1/source-depth-3m-12m-v5/

[R13] H10N MAERSK DETROIT exact-event direction. Quality-gated indeterminate Lane A result. Local: analysis/h10n-maersk-exact-event/v1/

[R14] H01W high-cadence identified-AIS replication design. Decision-ready only; no contact or acquisition authorized. Local: analysis/h01w-high-cadence-replication-design/v1/

External scientific sources retained by the project

[S1] Heaney et al. (2024), Regional soundscape modeling of the Atlantic Outer Continental Shelf, JASA 156, 378-390. DOI / source

[S2] MacGillivray and de Jong (2021), A reference spectrum model for estimating source levels of marine shipping based on AIS data, Journal of Marine Science and Engineering 9, 369. DOI / source

[S3] Australian Maritime Safety Authority monthly vessel traffic data and public release conditions. DOI / source

[S4] GEBCO gridded bathymetry data products. DOI / source

[S5] Copernicus Marine Global Ocean Physics Analysis and Forecast product family. DOI / source

[S6] Truong and Rogers (2021), blue-whale occurrence near Cape Leeuwin. DOI / source

[S7] Aulich et al. (2022), fin-whale acoustic presence around Australia. DOI / source

[S8] Gavrilov et al. (2011), harmonic structure of pygmy blue-whale song. DOI / source