Asia Green Corridors: A Zero-Backend Browser Observatory for Port and Shipping Decarbonisation Decision Support in Asia

Ka Wah Philip Ng
PhD, FCPA (Aust. & HK), FCMA, CESGA
Independent Researcher, Hong Kong
ngkwphilip@gmail.com · ORCID: 0009-0009-3955-010X

Abstract—Green shipping corridors are increasingly central to maritime decarbonisation in Asia, yet quantitative evidence to inform corridor choices remains limited and often requires paid data services. This paper introduces Asia Green Corridors, a zero-backend browser observatory that integrates previously separate public datasets: satellite-AIS port activity for 787 ports in PortWatch's Asia-Pacific classification (including 654 with Asian country codes) and its documented disruption log; satellite-estimated emissions from 2,885 Asian international-shipping sources (Climate TRACE); 534 western North Pacific cyclone tracks since 2005 (NOAA IBTrACS); and a live sea-state feed (Open-Meteo). The 2025 emissions layer totals 266.7 Mt of CO_2-equivalent on a 100-year global-warming-potential basis. The observatory provides one-click port screening, compares six Middle East–Asia corridor hubs by vessel calls and satellite-estimated emissions, and monitors three strategic chokepoints daily. A transparent scenario sandbox and checkpoint solver estimate the green-fuel shares implied by the IMO Revised Strategy's emissions checkpoints under an illustrative 2025-normalised pathway. At the model's default settings, meeting the indicative 2030 checkpoints on time requires deployed green-fuel shares of approximately 18%–29% of energy. These exceed the IMO's separate 5%–10% fuel-uptake ambition, but differing indicators, accounting assumptions and the model's restricted mitigation levers preclude interpreting the comparison as a direct policy shortfall. Anomaly detectors are evaluated against 83 eligible documented disruption events. At an operating point chosen by a transparent, predefined rule, they detected 59.0% of events in sample, generating 1.41 alerts per monitored port per year (2.04 per actively graded port); this event-level sensitivity is not a measure of comprehensive detection accuracy. The full sensitivity grid, event-type results and limitations are disclosed. Implemented as one HTML file and five static data bundles, the application requires no backend, account or API key and explicitly identifies unavailable feeds. It is openly released as an empirical companion to the author's corridor working paper.

Index Terms—*maritime decarbonisation, green shipping corridors, port analytics, automatic identification system (AIS), satellite-estimated emissions, anomaly detection, IMO strategy, Asia, open data, decision support, zero-backend web application.*

I. Introduction

A financier weighing a green-methanol bunkering terminal in Singapore, an insurer pricing war-and-weather risk for a Gulf–Asia tanker route, or a policy maker sizing a city's green-fuel ambition all face the same underlying question: what does the corridor actually look like — its traffic, its emissions, its vulnerabilities — and what would decarbonisation levers have to deliver? Answering it today typically means paid terminals or substantial manual data assembly. Yet the raw ingredients are public: satellite-AIS port activity, satellite-estimated shipping emissions, documented disruption events, and decades of cyclone tracks. What is missing is the join — a tool that brings these feeds together at the level of an individual port or corridor and states the result in plain language.

Asia Green Corridors is that join. It is a static, browser-executed application — one HTML file plus five static data bundles, requiring no application backend — that maps 787 ports in PortWatch's Asia-Pacific classification measured from satellite AIS, overlays satellite-estimated shipping emissions and cyclone exposure, tracks daily transits through the Malacca Strait, the Strait of Hormuz and the Taiwan Strait, and profiles six hubs of the Middle East–Asia green marine fuel corridor — Fujairah, Singapore, Hong Kong, Shanghai, Ulsan and Salalah — on measured annual vessel calls, vessel mix and emissions. A scenario sandbox models three decarbonisation levers against the IMO Revised Strategy checkpoints, and a checkpoint solver inverts the model to state, at the user's own settings, the green-fuel share of energy each checkpoint requires. The observatory was built as an empirical companion to my working paper on positioning Hong Kong in the Middle East–Asia green marine fuel corridor [6], whose limitations section notes that the corridor's quantitative evidence is "thin by construction"; this tool supplies an open, empirical data layer.

The entire system is zero-backend: one HTML file plus small static data bundles, fetching public feeds that permit cross-origin resource sharing (CORS) directly in the viewer's browser. There is no server, database, account, or API key; deployment is a drag-and-drop onto a static host. This paper describes the data estate (Section II), the architecture (Section III), the corridor analytics including the anomaly detectors and the scenario model (Section IV), a validation of the detectors against documented disruption events (Section V), a worked Hong Kong case study (Section VI), practitioner lessons (Section VII), and limitations (Section VIII).

The contributions are threefold: (1) an integrated, browser-only observatory joining port activity, emissions, disruption and cyclone datasets for PortWatch's Asia-Pacific port group and the Middle East–Asia corridor; (2) an anomaly-detection layer evaluated against documented disruptions, whose operating point is selected using a transparent, predefined rule and whose sensitivity grid, per-event-type performance and failure modes are published rather than hidden; and (3) a transparent, fully disclosed scenario model with an inverted checkpoint solver that converts IMO targets into required green-fuel shares. I claim no algorithmic novelty: the detectors and the scenario model are deliberately standard, and the contribution lies in the integration, the measured validation, and the accessibility.

II. Data Sources

The observatory joins four core public datasets plus a live sea-state feed (Table I); all are keyless and permit reuse under their respective terms, and weekly-updated sources carry their own "as of" dates in the interface. The join is the point: PortWatch supplies who calls where and when, Climate TRACE what that activity emits, the disruption log what has historically gone wrong, and IBTrACS the cyclone exposure each port lives with. Two provenance points matter for everything downstream. First, the Climate TRACE figures are the 2025 annual estimates returned by the Climate TRACE public API on 5 September 2026 — the API labels the gas basis co2e_100yr (a 100-year GWP basis) but does not state the assessment-report vintage, and returns no per-record release string, so the vintage of record is the access date; annual values remain subject to revision in later releases. Second, the annual vessel-call figures are as published in the PortWatch ports database, harvested on 5 September 2026; PortWatch does not state the exact 12-month window behind these published annual estimates. Geographic inclusion is explicit: the port universe is every port in PortWatch's own "Asia & Pacific" continent group — 787 ports, of which 654 carry Asian ISO3 country codes and the remainder are Pacific-island, Australian and New Zealand ports under that label — and the emissions subset is the Climate TRACE international-shipping sources whose ISO3 codes fall in a 46-code Asian list stated in the harvest script.

Table I. Data sources at a glance.

Source	Content	Cadence / vintage	Licence
IMF PortWatch ports & activity [1]	787 Asia-Pacific ports, vessel mix, daily calls, chokepoint transits	Weekly; series 2019–2026	IMF open data (cite PortWatch)
IMF PortWatch disruptions [2]	132 documented disruption records (validation reference set)	2018–2026	IMF open data
Climate TRACE [3]	2,885 Asian international-shipping sources, 266.7 Mt CO_2e (100-yr GWP)	2025 annual estimates; public API accessed 5 Sep. 2026	CC BY 4.0
NOAA IBTrACS [4]	534 western North Pacific cyclone tracks ≥34 kt	2005–2026, v04r01	Public domain
Open-Meteo Marine [5]	Sea-state context on port selection	Live	CC BY 4.0

Sources are pinged live from the viewer's browser on demand; a failed feed produces an explicit "unavailable" state, never a fabricated number.

III. Zero-Backend Architecture

The system is one HTML file (about 900 lines) plus five small static JavaScript data bundles harvested at build time: the port universe, the shipping-emissions sources, the disruption log, the cyclone tracks, and six hubs' full daily port-call series (2019–2026). Everything else — the map, charts, anomaly detectors, scenario model and solver — runs client-side in vanilla JavaScript with Chart.js and Leaflet [15], [16]. Live fetches are limited

to per-port daily activity (on selection), the chokepoint series, sea state, and the on-demand source-status check; each has an explicit grey "unavailable" degraded state, so the page remains fully usable on networks that block a feed. Data harvesting is a set of short Python scripts kept with the source; the browser-side anomaly detectors were cross-checked against the offline Python implementations on all six hubs and produce identical detections, so the validation results in Section V apply to the shipped code.

Zero-backend is not dependency-free: the page relies on third-party content-delivery networks for the mapping and charting libraries [15], [16] and basemap tiles [17], and on the continued availability of the live public APIs. The archived static bundles (Data and Code Availability) are the fallback, and every live panel degrades explicitly rather than silently. This architecture is a deliberate choice for the intended audience: a financier or policy officer can open the tool without local installation, while the public data sources, processing scripts and analytical methods can be independently inspected. The bundled inputs and scripts support reproduction of the figures in this paper.

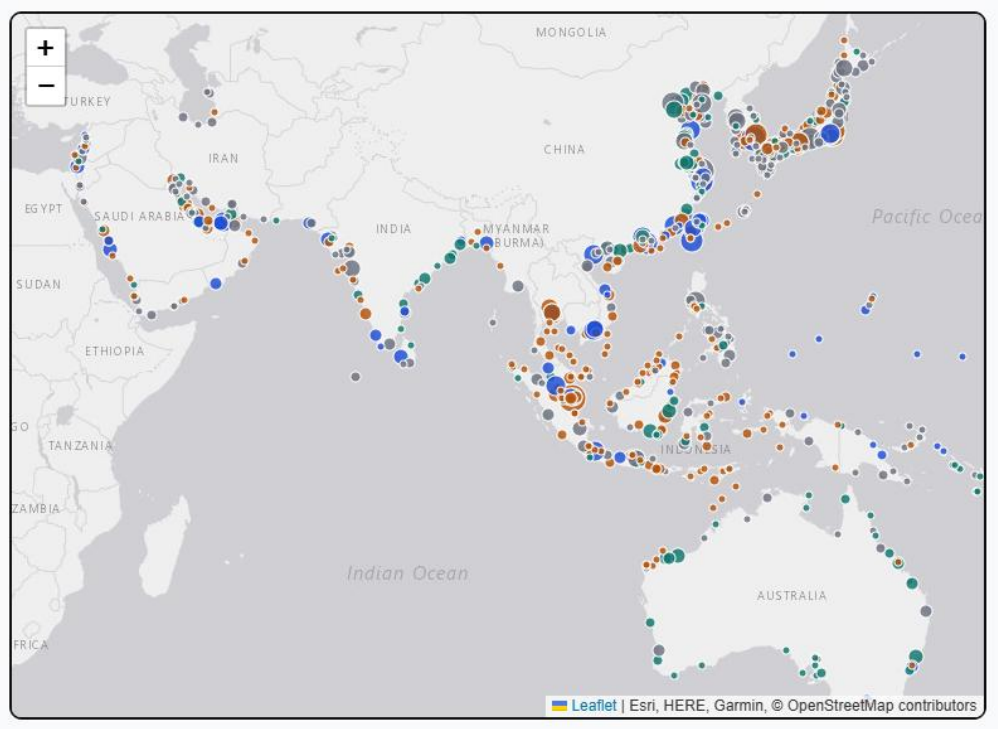


Fig. 1. PortWatch's Asia-Pacific port universe: 787 ports measured from satellite AIS (IMF PortWatch). Marker size is annual vessel calls; colour is dominant vessel type. Captured 2026-11-03.

IV. Corridor Analytics

A. Corridor hubs, measured

The working paper [6] compares Fujairah, Singapore and Hong Kong qualitatively; the observatory extends that comparison to six corridor hubs on measured annual vessel calls (PortWatch) and satellite-estimated shipping emissions (Climate TRACE, 2025) (Fig. 3). Singapore dominates on physical scale (44,369 calls, 21.58 Mt CO_2e); Fujairah is the tanker-intense midstream incumbent (91% tanker share); Hong Kong is a smaller physical node whose case, per the working paper's thesis, rests on institutional positioning rather than volume; Salalah appears with no Climate TRACE shipping source

— an explicit data gap, shown rather than filled. Two limitations are stated beside the table: AIS cannot see fuel type or cargo content, so tanker share is a proxy for liquid-bulk intensity, and green-capable vessels are not separately identifiable in the public feeds used here. The cyclone layer is likewise analytical: selecting a port counts the IBTrACS tracks that have passed within 300 km of it since 2005. Because PortWatch does not state the window behind its published annual figures, the bundled daily series allow two independent checks. First, a within-series year-over-year comparison needs no assumption about that window: summing each hub's trailing twelve months to 23 October 2026 against the preceding twelve months gives Fujairah 2,417 calls against 3,817 (−36.7%), Singapore 45,247 against 45,892 (−1.4%), Hong Kong 14,810 against 15,574 (−4.9%), Shanghai 36,613 against 43,488 (−15.8%), Ulsan 9,989 against 8,576 (+16.5%) and Salalah 1,730 against 1,604 (+7.9%). Fujairah's within-series drop — the largest decline among the six hubs, and straddling the disruption to Hormuz transits — coincides with that disruption and is consistent with substantial regional effects on shipping activity; a before-and-after comparison alone does not isolate its causal contribution, since rerouting, vessel-composition shifts and wider economic factors may also play a part. The hubs without direct Gulf exposure move between −15.8% and +16.5%, with Ulsan, the least Gulf-exposed, actually rising. Moreover, with roughly eight months of severely reduced transits inside that trailing window, the observed −37% decline and continued vessel calls indicate that Fujairah's activity was not entirely interrupted — consistent with the pipeline-fed resilience discussed below. Second, comparing the same trailing-twelve-month sums against PortWatch's published annual figures gives Fujairah −30% (2,417 summed versus 3,462 published), Singapore +2% (45,247 versus 44,369), Hong Kong −9% (14,810 versus 16,306), Shanghai −11% (36,613 versus 41,093), Ulsan −15% (9,989 versus 11,795) and Salalah +5% (1,730 versus 1,645). Ulsan's comparable shortfall without Gulf exposure means those gaps reflect methodological differences between the published annual estimates and the daily series as much as any crisis effect — the published-versus-summed comparison bounds the window ambiguity rather than isolating the disruption's contribution. The mixed direction makes the same point: the published figures for Ulsan and Hong Kong exceed both trailing twelve-month windows, while Singapore's sits below both. The published annual figures should therefore be treated as estimates with an unspecified reporting window, while the daily series provide a more recent signal for the Gulf hubs.

Daily chokepoint transits (Fig. 2) give the corridor a measured pulse — and, since early 2026, a demonstrated stress test. Hormuz transits collapsed from around 100 per day to near zero after 28 February 2026 and remained depressed through the northern summer. This bears directly on the corridor premise, and I flag it rather than leave it as a chart footnote: the 2025 emissions baseline is pre-crisis; the annualised hub volumes are PortWatch's current published figures, with an unspecified reporting window (see the trailing-twelve-month comparison above); the corridor's route risk is now demonstrated rather than hypothetical; and Fujairah's measured tanker intensity should be read in the light of a port that sits outside the strait and can be fed by pipeline from Abu Dhabi when transits through the strait are severely disrupted. A corridor whose Gulf end depends on a chokepoint whose transits have now been severely disrupted for months is a corridor whose Asian end matters more, not less.

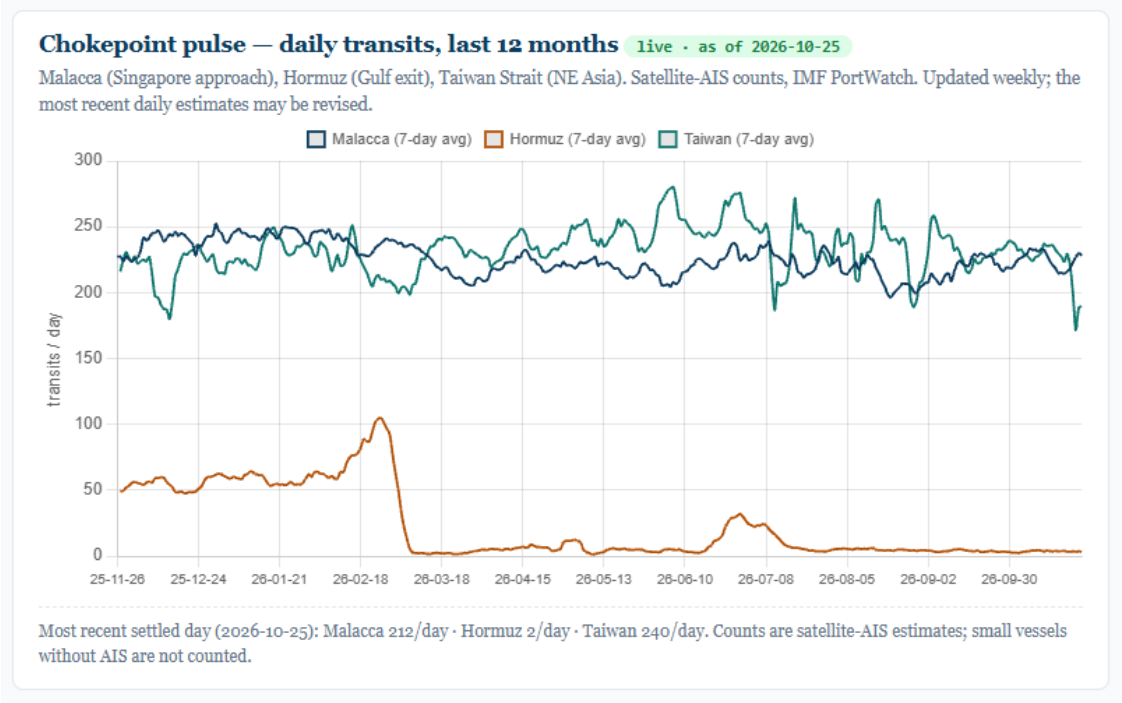


Fig. 2. Chokepoint pulse: seven-day-average daily transits through Malacca, Hormuz and the Taiwan Strait, last twelve months (IMF PortWatch, captured 2026-11-02, data as of 2026-10-25). The collapse in Hormuz transits from late February 2026 reflects the 2026 Gulf crisis.

Corridor hubs, measured — extending the three-hub comparison

Ng (2026) compares Fujairah, Singapore and Hong Kong qualitatively. Here, that comparison is extended to six corridor hubs using measured annual vessel calls (PortWatch) and satellite-estimated shipping emissions (Climate TRACE, 2025).

Hub	Vessel calls / yr	Tanker share	Shipping emissions (Mt CO_2e)	Activity (M nmi)	Reading
Fujairah	3,462	91%	1.3	4	Incumbent midstream — 9 of 10 calls are tankers (Ng 2026: "capacity-led")
Singapore	44,369	59%	21.58	46.8	Full-stack incumbent — largest physical node in Asia
Hong Kong	16,306	12%	3.5	9.2	Smaller physical node; institutional positioning (Ng 2026 thesis)
Shanghai	41,093	20%	8.15	20.9	Largest Mainland gateway; CT port-area cluster
Ulsan	11,795	73%	2.32	9.2	Korea's designated green ship-fuel port
Salalah	1,645	11%	—	—	Omani green-ammonia base — not yet a CT shipping source (honest gap)

Tanker share is a proxy for fuel/oil handling intensity — AIS cannot see fuel type or cargo content, so a tanker call is evidence of liquid-bulk activity, not of green-fuel trade. Green-capable vessels are not separately identifiable in the public feeds used here; this is a data limitation, not a classification or rounding choice.

Fig. 3. Six corridor hubs measured: annual vessel calls and tanker share (PortWatch ports database, harvested 2026-09-05; the exact 12-month window behind PortWatch's published annual figures is

not stated) and shipping emissions (Climate TRACE 2025 annual estimates), with the limitation note shown to users; see the trailing-twelve-month comparison in Section IV-A.

B. Anomaly detection

Two complementary detectors run on each hub's full daily series. A structural detector (binary-segmented Pettitt change-points [9], $p < 0.05$, minimum segment 90 days, ≥25% level shift, at most 8 splits) finds persistent regime changes such as closures, conflict-driven rerouting, or capacity shifts; binary segmentation applies the test repeatedly without a multiple-testing correction, so structural detections should be read as screening flags. A transient detector flags sudden dips — the seven-day trailing mean falling below a set fraction of the trailing 90-day median for at least three days, closing at ≥80% recovery — suited to short shocks such as cyclones and floods. Detections are matched against PortWatch's documented disruption log within ±14 days, and unmatched alerts are labelled "No match in PortWatch DB" rather than dismissed: an unmatched alert is not necessarily a false alarm, because the log is not exhaustive. At the adopted setting the six hubs show 6 structural change-points and 17 transient alerts since 2019; the interface displays the most recent detections with their match status.

C. Scenario sandbox and checkpoint solver

The sandbox models three levers against the satellite-estimated 2025 baseline: green-fuel share of energy by 2035, fleet share adopting slow steaming (power demand varies roughly with the cube of speed, but a slower voyage takes longer, so fuel per voyage varies roughly with its square — a 10% speed reduction cuts fuel per voyage by roughly 19%), and the share of port calls on shore power (removing roughly 4% of voyage emissions at berth). Levers phase in linearly 2026–2035 and hold. One-for-one substitution assumes the green fuel is zero-carbon on a well-to-wake basis; e-methanol, bio-methanol and biofuels have non-zero well-to-wake intensity, and a residual intensity ρ would inflate the required green-fuel share by a factor of $1/(1-\rho)$. The IMO Revised Strategy checkpoints (2030: −20% striving for −30%; 2040: −70% striving for −80%; net zero by or around 2050 — the reductions defined relative to 2008 [7], whose baseline level is documented in the Fourth IMO GHG Study [10]) are overlaid on the 2025 baseline as an approximation. Two boundary caveats are printed beside the model: the checkpoints are defined in the IMO Strategy relative to 2008, so if 2025 Asian shipping emissions exceed their 2008 level the true required reductions would be greater than those shown; and Climate TRACE emissions are attributed to port-area clusters, not the IMO's global international-shipping accounting boundary, so the overlay is a proportional analogue rather than a subset of the IMO baseline. The model is labelled a heuristic planning tool, not a forecast.

The checkpoint solver inverts the model: at the user's current slow-steaming and shore-power settings it reports the model-implied green-fuel share required for the modelled path to reach each checkpoint (Table II). The outputs are model-implied fuel shares under an illustrative 2025-normalised emissions pathway — not estimates of fuel demand, verified lifecycle emissions reductions, or official port-specific compliance obligations. The arithmetic is fully explicit: the modelled reduction phases in linearly and reaches its full value in 2035 — $r(t) = f(t)\cdot[1 - (1 - g)K]$, with $f(t) = (t - 2025)/10$ clamped to $[0, 1]$ and evaluated at the checkpoint year (so $f(2030) = 0.5$), and $K = (1 - 0.19s)(1 - 0.04p)$ at slow-steaming adoption s and shore-power share p. Inverting, $g = 1 - (1 -$

r/f(t))/K — every figure in Table II is recomputable by hand at the sandbox defaults (s = 30%, p = 20%). Table II reports terminal (2035) green-fuel shares; the deployed share in the checkpoint year is g·f(t), so the 2030 rows correspond to 2030 deployed shares of 17.9% (minimum) and 28.6% (striving), and for r > f(t) no green share suffices (g > 1) — the feasibility boundary the 2050 row sits on. Because all levers are proportional, these percentages are identical across scopes — all Asian international shipping, corridor hubs only, or Hong Kong alone; only the absolute emissions amounts change. This scope invariance is a property of the model, not of the world: applying a global checkpoint rate to a single port or corridor implicitly assumes equal-rate burden sharing, which the IMO Strategy does not prescribe. The solver also evaluates the 2035 plateau against each checkpoint's level, so reaching the 2030 checkpoint at the plateau means reaching it five years late; meeting it in 2030 would require materially more, and the 2050 net-zero goal requires a 100% green share under one-for-one substitution — i.e. the model offers no path to net zero short of complete fuel substitution, and the efficiency levers cannot close any part of the final gap. The comparison with the Strategy's separate fuel-uptake ambition is illustrative: the Revised Strategy calls for zero or near-zero GHG emission fuels, technologies and/or energy sources to represent at least 5%, striving for 10%, of the energy used by international shipping by 2030 [7] — yet under the model's simplified assumptions, meeting the indicative 2030 emissions checkpoints on time would require deployed green-fuel shares of about 18%–29% of energy by 2030 (roughly 3.6 times the 5% minimum ambition at the minimum checkpoint, and 2.9 times the 10% striving ambition at the striving checkpoint). The two quantities concern different indicators within a broader framework, and the sandbox uses restricted levers and a 2025-normalised baseline, so the difference indicates scale rather than a directly comparable policy shortfall. Within this restricted three-lever model, the implied fuel-substitution requirement remains substantially greater than the separate IMO fuel-uptake ambition, illustrating the scale of the decarbonisation challenge under the model's assumptions. As of writing, the IMO Net-Zero Framework intended to give the Strategy legal force remains unadopted: the October 2025 extraordinary session adjourned without an adoption vote, MEPC 84 (April–May 2026) reached no agreement, and the resumed extraordinary session scheduled for 4 December 2026 remains subject to confirmation [8]. The Framework's GHG fuel-intensity metric maps directly onto the green-fuel-share lever modelled here.

Table II. Checkpoint solver output at the default settings (30% slow-steaming adoption, 20% shore power): terminal (2035) green-fuel share of energy required. Deployment in the checkpoint year is the terminal share × f(t) — half at 2030 (f = 0.5). Percentages are scope-invariant under proportional levers; absolute emissions scale with scope.

IMO checkpoint	Terminal (2035) share — checkpoint met in its year	Terminal (2035) share — checkpoint met at the 2035 plateau
2030 minimum (−20%)	35.9%	14.5%
2030 striving (−30%)	57.2%	25.2%
2040 minimum (−70%)	67.9%	67.9%

IMO checkpoint	Terminal (2035) share — checkpoint met in its year	Terminal (2035) share — checkpoint met at the 2035 plateau
2040 striving (−80%)	78.6%	78.6%
2050 net zero (by or around 2050)	100%	100% — no path short of complete fuel substitution

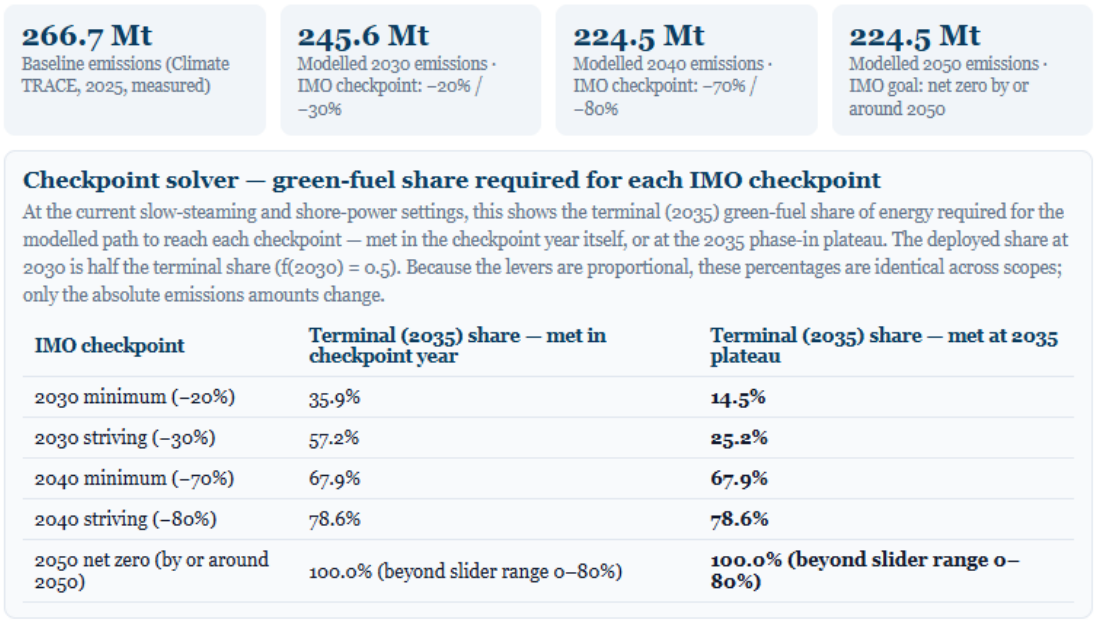


Fig. 4. Scenario sandbox outputs and the checkpoint solver (values in Table II): modelled 2030/2040/2050 emissions against the IMO Revised Strategy checkpoints; baseline 266.7 Mt CO_2e (Climate TRACE 2025 annual estimates). Captured 2026-11-02.

V. Validation against Documented Disruptions

Because the observatory draws conclusions from satellite ship-tracking data, its alarm system is tested against documented reality and the results published, including where the method performs poorly. The validation uses two public PortWatch datasets: daily port traffic for the 283 ports worldwide for which a complete daily series is published over 2019–2026, and the documented disruption log. (The tool's 787-port Asia-Pacific map includes ports whose daily series is not published in full; the validation monitors every port worldwide with a complete published series, which is why the monitored universe extends beyond Asia.) The bundled log contains 132 disruption records; 130 have usable event dates, and of these 83 are eligible for grading. The 47 excluded records break down as follows: 45 name no specific monitored port (country- or basin-level entries such as "Earthquake in Peru"), and the remaining two pre-date the baseline window — an event is graded only when at least 60 days of pre-event traffic exist at an affected port (the transient detector's baseline median is computed on a trailing 90-day window, so it can fire no earlier than day 91 of a series; the only eligible events with 60–89 days of pre-event traffic — Cyclone Idai and the contemporaneous Mozambique flood, both March 2019 — could therefore be caught only by the structural detector, and both are graded misses). Events the data cannot see are never tested. An event is counted as

caught if an alert fires at an affected port within ±14 days of the documented dates. The sensitivity grid varies only the transient detector; the structural detector runs unchanged in the background and contributes materially — of the 49 events caught at the adopted setting, 4 were matched only by the structural detector and 15 by both detectors — but because it is invariant across the grid, the graded trade-off remains the transient detector's.

The operating point is selected using a transparent, predefined rule on the published six-row sensitivity grid (Table III) — the highest catch rate subject to no more than two alerts per port per year. Because the rule is applied to the same events used for grading, this remains an in-sample selection; an independent test period would strengthen future validation. At the adopted setting (transient detector at 60% of normal, grading ports busier than one call per day), the system caught 49 of 83 documented events — 59.0%, an in-sample, event-level sensitivity against the eligible reference set rather than comprehensive detection accuracy; the 83 events are a screening reference set, not a representative sample of all maritime disruptions. (Excluding the two structurally-only-testable events would give 49 of 81, or 60.5% — keeping them as misses is a conservative bias of about 1.5 percentage points.) Across the full observation period (1 January 2019 – 23 October 2026, 7.81 years) it generated 3,125 alerts across all 283 monitored ports — 1.41 alerts per port per year. Of the 3,109 alerts at event-affected ports (the 279 monitored ports named in at least one eligible event, which is why nearly all alerts occur there), 130 coincided with a documented event — 4.2%. (The four monitored ports never named in an eligible event account for the remaining 16 alerts of the 3,125 — 0.51 per port per year.) That 4.2% is an observed event-match rate, not an estimate of detection precision, because the reference log is incomplete and temporal matches do not independently establish causality: COVID-19 congestion and the 2026 Gulf crisis, both plainly visible in the traffic data, are not catalogued as events, and since the disruption log derives largely from natural-hazard reporting, the true precision cannot be estimated from this reference set.

Table III. Measured validation results: sensitivity grid (tolerance ±14 days; 83 eligible events).

Sensitivity setting	Total alerts (all 283 ports)	Ports graded by filter	Disruptions caught	Alerts / port / year (of 283)	Alerts matching a documented event (observed match rate)
60% of normal × 1.0 call/day (adopted)	**3,125**	**196**	**59.0% (49/83)**	1.41	130 of 3,109 (4.2%)
60% of normal × 3.0 calls/day	1,433	113	42.2% (35/83)	0.65	66 of 1,417 (4.7%)
50% of normal × 1.0 call/day †	1,912	196	43.4% (36/83)	0.87	80 of 1,898 (4.2%)
50% of normal × 3.0 calls/day, ≥4 days	1,108	113	33.7% (28/83)	0.50	38 of 1,095 (3.5%)

Sensitivity setting	Total alerts (all 283 ports)	Ports graded by filter	Disruptions caught	Alerts / port / year (of 283)	Alerts matching a documented event (observed match rate)
70% of normal × 1.0 call/day	4,886	196	68.7% (57/83)	2.21	195 of 4,865 (4.0%)
70% of normal × 0.5 calls/day, ≥2 days	9,073	231	75.9% (63/83)	4.11	325 of 9,041 (3.6%)

Denominators are held at 83 across rows for comparability; events at ports excluded by a row's activity filter are counted as misses, not abstentions. Total alerts are across the full monitored universe, and the alerts-per-port-per-year column divides by all 283 monitored ports for comparability. "Ports graded" counts the ports whose trailing 90-day median traffic reaches the row's activity floor at least once over 2019–2026; the burden per watched port is therefore higher than the headline rate — at the adopted setting, 3,125 alerts across 196 graded ports is 2.04 per watched port per year, and the declared ≤2 rule is defined per monitored-universe port. Match counts are against alerts at event-affected ports. The adopted row is shown in bold; † marks the setting the ≤2 rule would select under a watched-port denominator (see text).

Three further findings matter for interpretation. First, performance varies by event type in a way that is consistent with what the detector measures: at the adopted setting it caught 38 of 62 cyclones and 6 of 8 floods, but only 3 of 10 earthquakes — the detector sees disruption to vessel activity, not physical damage to port infrastructure, and earthquakes often leave shipping largely undisturbed. (The remaining three of the 83 events — two conflict/other and one volcanic eruption — were caught twice and missed once.) Second, the result is not driven by the matching window: widening the tolerance from ±7 to ±21 days moves the catch rate only from 51.8% to 62.7%. Third, restricting the analysis to the Asian subset yields 30 of 54 events detected (55.6%), broadly comparable with the global in-sample result, although independent validation would be needed to establish generalisability. Full event-level results ship with the code (validation/results.json).

For the intended early-warning use, the decision-relevant summary is the trade-off across Table III: catching more than half of documented disruptions costs 1.41 alerts per monitored port per year — or 2.04 per watched (graded) port, the figure an insurer or operator actually staffs against; pushing sensitivity to catch three-quarters costs 4.11 per monitored port, or 5.03 per watched port. One denominator-sensitivity disclosure belongs here: under a watched-port denominator the adopted row's 2.04 alerts per year would fail the declared ≤2 rule, which would instead select the 50% × 1.0 call/day setting (43.4% caught, 1.25 per watched port per year; row marked † in Table III). I retain the monitored-universe denominator for comparability across rows and report both burdens. An insurer and a port operator may reasonably choose different settings, and the interface reports the whole grid rather than a single flattering number. My disclosure principle throughout: where the method is weak, the number is printed, not hidden.

Sensitivity setting (alert when traffic falls below … × only grade ports busier than …)	Disruptions caught (of 83)	Alerts per port per year	Alerts matching a documented event (observed match rate)
60% of normal × 1.0 call/day **(adopted)**	**59.0%** (49/83)	1.41	4.2%
60% of normal × 3.0 calls/day	42.2% (35/83)	0.65	4.7%
50% of normal × 1.0 call/day	43.4% (36/83)	0.87	4.2%
50% of normal × 3.0 calls/day, ≥4 days	33.7% (28/83)	0.50	3.5%
70% of normal × 1.0 call/day	68.7% (57/83)	2.21	4.0%
70% of normal × 0.5 calls/day, ≥2 days	75.9% (63/83)	4.11	3.6%

How to read a row: **"60% of normal × 1.0 call/day"** means — raise an alert when a port's 7-day average traffic falls below 60% of *its own* recent normal (the trailing 90-day median) for at least 3 days; and only grade ports whose normal traffic is at least 1.0 call/day, because quieter ports are statistically too noisy to judge. All columns are defined in steps III–IV above.

Fig. 5. The self-grading panel as shown to users (values in Table III): the sensitivity grid with the adopted operating point, and the plain-language "how to read a row" note. Captured 2026-11-03.

VI. Case Study: Hong Kong in the Corridor

Hong Kong illustrates how the observatory's layers combine for a policy question. The measured hub comparison shows a smaller physical node: 16,306 calls a year — about a third of Singapore's 44,369 — and 3.5 Mt CO_2e of port-area shipping emissions, roughly a sixth of Singapore's 21.58 Mt, with a 12% tanker share. The scenario solver states the decarbonisation requirement in terms a policy maker cannot avoid: on the model's implicit equal-rate burden-sharing assumption, Hong Kong faces the same arithmetic as the corridor — a terminal green-fuel share of about 25% of energy by 2035 to reach the 2030 striving level five years late — or 57% terminal (28.6% deployed by 2030) to meet it on time — and about 79% for the 2040 striving checkpoint, at the default settings for the other levers, which contribute little by comparison. These are modelled requirements against an approximate (2025-based) overlay of Strategy checkpoints defined relative to 2008; they represent substantial illustrative requirements rather than verified compliance obligations or feasibility judgements. For calibration, Hong Kong's Action Plan on Green Maritime Fuel Bunkering projects over 200,000 tonnes of green fuel bunkered a year and at least 60 bunkering operations by 2030, and targets at least an 11% reduction in carbon emissions from Hong Kong-registered ships by 2026 against 2019, driven by a Carbon Intensity Indicator incentive scheme [11], [13]. The scheme's application window runs to 31 December 2026 [13] (1,509 approvals covering 974 ships as at 28 February 2026 [18]), so whether the 11% reduction was met is not yet reportable. The Action Plan's headline 7% figure is officially described in two different ways: the Action Plan and a 2026 Legislative Council paper frame it as 7% of Hong Kong-registered ships taking up green maritime fuels by 2030 — a fleet-uptake share [11], [13] — while 2025 Legislative Council papers frame it as green maritime fuels accounting for 7% of all fuels of Hong Kong-flagged vessels — an energy share [12]. The denominators differ and are worth stating: the 200,000-tonne figure is the volume bunkered in Hong Kong, whereas the 7% energy-share reading concerns Hong Kong-flagged vessels wherever

they fuel. Both figures are also more permissive than the solver's lever: Hong Kong's "green maritime fuel" category includes LNG and biodiesel, which carry non-zero well-to-wake intensity, whereas the modelled lever assumes zero-carbon substitution — further limiting direct comparison with the modelled 18%–29% range. (Over 220,000 tonnes were bunkered in 2025, mostly biodiesel — already above the Action Plan's 200,000-tonne figure, though the Plan frames that volume around LNG- and methanol-powered vessels, for which only 16 LNG bunkering operations took place — so the 2025 outturn does not by itself establish fulfilment of the Plan's 2030 objective [13].) On the second reading the target is the closest available comparison to the solver's output, and 7% by 2030 sits well below the roughly 18–29% 2030 deployed green share the 2030 checkpoints imply on this model. Recent policy developments reinforce the institutional reading: the 2026 Policy Address announces a "volume-to-value" transformation of the port towards higher-value maritime services, ship-registration and marine-insurance measures, and Hong Kong's first "Green Energy Corridor" with a Mainland city [19].

Read together with the working paper [6], the empirical evidence supports an institutional-platform role for Hong Kong, emphasising standards, finance, verification and facilitation as potential complements to the greater physical scale of hubs such as Singapore. For Asian policy makers more broadly, the same solver arithmetic applies to every scope, and the hub table quantifies where the physical volumes actually sit; for the working paper, the observatory supplies the open, reproducible evidence base it explicitly lacked.

VII. Lessons Learned

1) Weekly feeds need "as of" dates. Port activity and chokepoint feeds are updated weekly with revisable recent days; every panel carries its own date stamp so a stale number is never mistaken for a live one.

2) Treat reference sets as incomplete by design. A documented-event log built for situational awareness is not an exhaustive catalogue of disruptions. Scoring against it yields an observed event-match rate, not precision, and the validation text says so before the reader can misread 4.2% as "96% false alarms".

3) Quiet ports are difficult to grade reliably. At the adopted setting, ports below about one call per day are treated as statistically too noisy, and the interface shows an explicit "not enough data" state rather than guessing. The sensitivity grid is stricter: it holds the denominator at 83 (Table III note), so events at filtered-out ports are counted as misses and depress the reported catch rate — the published figures are conservative in that direction.

4) Reproducible arithmetic beats impressive numbers. Publishing the exact numerator, denominator and period for every rate ($3{,}125 \div 283 \div 7.81 \approx 1.41$) lets an analytically minded reader recompute every headline figure — and during preparation this discipline caught a stale rate left over from an earlier harvest.

5) Cross-check the browser against the reference implementation. The shipped JavaScript detectors were verified against the offline Python versions on all six hubs; without that step, the validation numbers would describe different code than the one users run.

6) Explicit degraded states preserve trust. Blocked feeds produce grey "unavailable" states, never interpolated values; throughout testing, missing data was shown as data.

VIII. Limitations and Future Work

Intended use — Asia Green Corridors is a decision-support screening tool for project financiers, investors, insurers and policy makers. It is not a trading system, not a compliance calculation, and not a substitute for verified IMO Data Collection System (DCS) or EU Monitoring, Reporting and Verification (MRV) data [14]. Nothing in it constitutes investment, insurance or policy advice.

AIS activity cannot identify fuel type or cargo, so tanker share is only a proxy for liquid-bulk intensity, and green-capable vessels are not identifiable in the public feeds used. AIS coverage misses some small vessels. Emissions are Climate TRACE satellite estimates for 2025 (a release subject to revision), not measured fuel sales, and are attributed to port-area clusters rather than individual berths; one corridor hub (Salalah) has no matching shipping source and is shown as an explicit gap. The scenario model is a transparent heuristic with deliberately simple assumptions — one-for-one zero-carbon fuel substitution, a speed–power cube-law and voyage-fuel square-law approximation, a fixed shore-power fraction, linear phase-in — and its checkpoint overlay is approximate because IMO targets are defined against 2008 while the satellite-estimated baseline is 2025, with a different accounting boundary. Slow steaming is modelled as a pure efficiency gain; in practice lower speeds require more ship-days — and eventually more vessels — for the same trade volume, partially offsetting the saving. The solver's outputs are model-implied fuel shares under an illustrative 2025-normalised emissions pathway — not estimates of fuel demand, verified lifecycle emissions reductions, or official port-specific compliance obligations. The validation is in-sample: it inherits the incompleteness of the event log, conservatively counts eligible events at ports excluded by the activity filter as misses, and cannot estimate true precision from the available reference set. Planned extensions include port-level emissions allocation, a wider set of corridor hubs, multi-year emissions vintages as Climate TRACE releases them, and propagation of input uncertainty into the scenario outputs.

IX. Conclusion

Asia Green Corridors demonstrates how publicly accessible maritime datasets can be integrated into a browser-based observatory to support transparent screening of Asia's shipping decarbonisation challenges. The tool combines port activity, satellite-estimated emissions, chokepoint monitoring and scenario analysis, and evaluates its anomaly-detection performance against documented disruptions — publishing the misses alongside the hits. Its principal contribution is an accessible and reproducible empirical framework rather than algorithmic novelty or compliance-grade forecasting. By making its assumptions, data gaps and methodological limitations explicit, the observatory provides a practical foundation for further corridor analysis, policy assessment and investment screening.

Acknowledgment

The author thanks the IMF PortWatch team, Climate TRACE, NOAA NCEI (IBTrACS), and Open-Meteo for maintaining open data services, and the OpenStreetMap, CARTO and Esri contributors for the basemap.

Disclosure of AI Assistance

This system was designed and implemented with AI assistance (Kimi AI) acting as a pair-programming and drafting tool under continuous human direction. All engineering decisions, data-validation steps, and the final text of this manuscript were reviewed, tested, and approved by the human author, who takes full responsibility for the content.

Data and Code Availability

Source code and data bundles are available under the MIT licence on GitHub (github.com/drngkwphilip/asia-green-corridors) and permanently archived on Zenodo (doi: 10.5281/zenodo.23920029); the live demonstration is hosted as a static site at asia-green-corridors.netlify.app. Reproducibility: Climate TRACE 2025 annual estimates returned by the public API, accessed 5 September 2026 (a raw API response, with HTTP headers, is archived with the code as ct_api_raw_sample.json); PortWatch ports and disruption log harvested 5 September 2026, daily port-call series refreshed 3 November 2026; IBTrACS v04r01; validation window 1 January 2019 – 23 October 2026; interface screenshots captured 6 September 2026 (Fig. 3, one day after the first data harvest), 2 November 2026 (Figs. 2 and 4) and 3 November 2026 (Figs. 1 and 5); harvest scripts and the validation harness (validation_harness.py, results.json) ship with the code. Data: IMF PortWatch [1], [2]; Climate TRACE [3]; NOAA IBTrACS [4]; Open-Meteo [5].

Disclaimer

The views and opinions expressed in this paper are solely those of the author and do not reflect the official policies or positions of any current or former employers, professional associations, or affiliated institutions.

This article is intended strictly for educational and informational purposes. It does not constitute financial, legal, tax, investment, or scientific advice, nor should it be construed as a recommendation for any specific commercial strategy. Given the complex and dynamic nature of the regulatory and macroeconomic environments discussed herein, no warranties are made regarding the ongoing accuracy or completeness of the information presented.

Readers are advised to conduct independent verification and consult qualified professionals prior to making any strategic decisions. The author disclaims any liability for actions taken in reliance upon the contents of this publication.

References

[1] IMF, "PortWatch: Port activity and chokepoint transits from satellite AIS," portwatch.imf.org, accessed Sep. 5, 2026.

[2] IMF, "PortWatch disruptions database, 2018–2026," portwatch.imf.org, accessed Sep. 5, 2026.

[3] Climate TRACE, "Emissions inventory and public API: 2025 annual estimates for international shipping (CO_2e, 100-yr GWP)," climatetrace.org, accessed Sep. 5, 2026.

[4] K. R. Knapp, M. C. Kruk, D. H. Levinson, H. J. Diamond, and C. J. Neumann, "The International Best Track Archive for Climate Stewardship (IBTrACS)," Bull. Amer. Meteor. Soc., vol. 91, no. 3, pp. 363–376, 2010, doi: 10.1175/2009BAMS2755.1.

[5] P. Zippenfenig, "Open-Meteo.com Weather API," Zenodo, 2023, doi: 10.5281/zenodo.7970649.

[6] K. W. P. Ng, "From Port to Platform: Positioning Hong Kong as the Facilitator of the Middle East–Asia Green Marine Fuel Corridor under the First Five-Year Plan (2026–2030)," working paper, Zenodo, 2026, doi: 10.5281/zenodo.21639017.

[7] IMO, "2023 IMO Strategy on Reduction of GHG Emissions from Ships," Resolution MEPC.377(80), Jul. 2023.

[8] IMO, "Outcomes of MEPC 83 (Apr. 2025), the second extraordinary session (Oct. 2025) and MEPC 84 (Apr.–May 2026) on the Net-Zero Framework," imo.org, accessed Sep. 5, 2026.

[9] A. N. Pettitt, "A non-parametric approach to the change-point problem," J. R. Stat. Soc. Ser. C, vol. 28, no. 2, pp. 126–135, 1979, doi: 10.2307/2346729.

[10] IMO, "Fourth IMO Greenhouse Gas Study 2020," International Maritime Organization, London, 2021.

[11] Government of the Hong Kong SAR, "Action Plan on Green Maritime Fuel Bunkering," Nov. 2024, and press release, Nov. 15, 2024.

[12] Hong Kong Legislative Council Panel on Economic Development, papers on the Action Plan on Green Maritime Fuel Bunkering, Feb. and Mar. 2025.

[13] Hong Kong Legislative Council Panel on Economic Development, papers on green maritime fuel bunkering and the Action Plan, papers CB(2)342/2026(04) and CB(2)342/2026(05), Mar. 2026.

[14] IMO, "Data Collection System for fuel oil consumption of ships," Resolution MEPC.278(70), 2016; EU, "Regulation (EU) 2015/757 on the monitoring, reporting and verification of carbon dioxide emissions from maritime transport (EU MRV)," 2015.

[15] Chart.js v4.4.3, chartjs.org, accessed Sep. 5, 2026.

[16] Leaflet, leafletjs.com, accessed Sep. 5, 2026.

[17] Basemap tiles © OpenStreetMap contributors, © CARTO, © Esri, accessed Sep. 5, 2026.

[18] Hong Kong Transport and Logistics Bureau, "Controlling Officer's Reply No. 3635: Marine Department, 2026–27 Estimates," 2026.

[19] The Chief Executive of the Hong Kong SAR, "The 2026 Policy Address," Sep. 2026; Hong Kong Maritime and Port Development Board, "HKMPDB welcomes Hong Kong's First Five-Year Plan and the 2026 Policy Address," press release, Sep. 16, 2026.