EcoRouter

Cloud-based weather routing for safe, energy-efficient and just-in-time vessel voyages.

Routing in a new dimension

EcoRouter is a SINTEF Ocean cloud platform that gives voyage planners a complete, physics-based picture of their next crossing – every Pareto-optimal trade-off between transit time and energy consumption, evaluated under real forecast or historical weather. The browser is the only client; the heavy work runs on dedicated Azure compute provisioned per user. The platform supports decision making at three planning horizons:

- **Strategic:** Testing new technologies before they reach production (wind-assisted propulsion, alternative hull shapes, novel power trains).
- **Tactical:** Fleet performance analysis: comparative analysis of departure times and the resulting impact on energy consumption and safety.
- **Operational:** Online weather routing for safe, energy-efficient routes with just-in-time arrival, based on live ECMWF / Copernicus forecasts.

Platform at a glance

Fig. 1: EcoRouter platform overview. Voyage plan, weather data, and ShipX-derived ship models flow into the route optimiser, returning a Pareto set of routes.

- Built on SINTEF Ocean's ShipX hydrodynamics
- Three case modes: Optimisation, Comparison, Analysis
- Reference-route corridors auto-generated, drawn manually, or imported from RTZ / CSV
- Per-user Azure container – 16 vCPU / 64 GB RAM, on-demand
- Live job tracking and weather playback along every route
- Multi-tenant by design (JWT auth, ADMIN/USER roles)

Cloud architecture

EcoRouter is built around a permanent API server on Azure (Spring Boot on App Service) that orchestrates everything. It handles authentication with short-lived JWT tokens, persists every voyage, vessel and result in PostgreSQL, reads ship-model and coastline files from Azure Files, and never performs heavy computation itself. All optimisation runs on a per-user worker – an Azure Container Instance (16 vCPU / 64 GB RAM) provisioned the moment the user logs in and torn down automatically after idle. The worker mounts the same Azure Files share, pre-fetches weather grids from external providers, runs the DP optimiser, and writes Pareto results back to PostgreSQL – all while streaming live status to the browser.

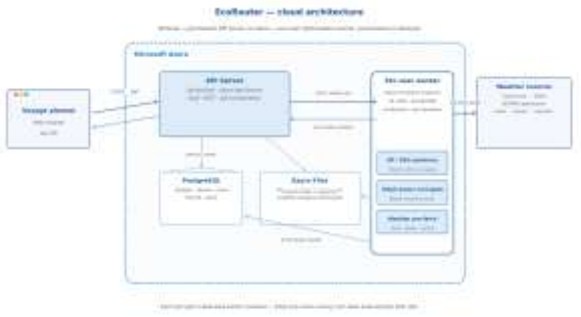

Fig. 2: EcoRouter cloud architecture (Azure deployment).

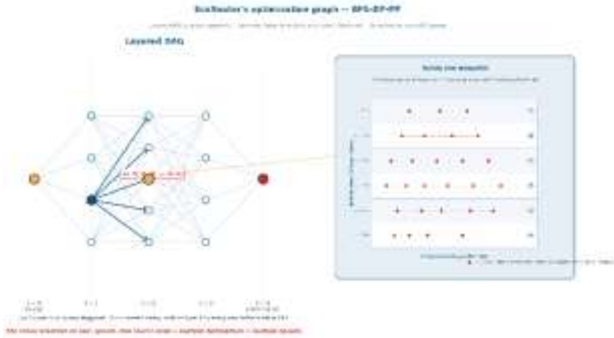

Fig. 3: Layered optimisation graph.

Optimisation engine

The optimiser is a Dynamic Programming (DP) algorithm over a layered Directed Acyclic Graph (DAG) of ocean waypoints. The DAG structure lets DP build optimal partial solutions layer by layer with no redundant recomputation. The objective is biobjective – transit time vs. energy consumption – and rather than collapse the two into a single weighted scalar, the solver keeps both alive throughout the search. Each node holds a *Pareto set* of non-dominated partial paths: every route reaching that waypoint that cannot be strictly beaten on both time and energy at once.

Branching is exhaustive: at every node the solver considers all outgoing arcs and the full discrete speed range (default 0.25 kt steps), calling the vessel power model with the actual sea state – wave height, period and direction, wind speed and direction, ocean current – interpolated from the downloaded weather grids. Operational safety is enforced at the same stage: legs whose forecast conditions exceed user-defined thresholds are pruned from the search before any cost is computed.

Dashboard

After logging in, the planner lands on the Dashboard. A status banner shows the live state of the user's personal optimisation engine – *Starting up*, *Ready* or *Stopped* – together with provisioning age and remaining idle time before automatic teardown. From the Dashboard the planner can navigate to Ports, Reference Routes, Cases, and (for administrators) Vessels, Users, World Map and worker management.

Fig. 4: Personal dashboard.

Case modes

Optimisation

The core mode. The planner works with a single vessel, selects one or more reference routes for a port pair and runs the full DP solver. The output is the complete Pareto front for that vessel on that crossing under forecast weather – a clear, data-driven menu of options to choose from.

Vessels Comparison

For fleet managers, naval architects and research teams who need to benchmark two ships on the same scenario. The case is configured identically to an Optimisation case except that two vessels are selected; the system runs the full optimisation independently for each and presents both Pareto fronts side by side.

Analysis

Generalises the single-departure run into a systematic sweep. The system steps through a user-defined range of departure offsets, solves each as a full optimisation and aggregates all results into a single departure-time sensitivity chart. This reveals the best weather window in the planning horizon (e.g. leaving 12 h later saves 8% fuel by letting a deep low pass before the vessel reaches the open ocean).

Fig. 5: Case-mode selection.

Ports

Ports are the building blocks of every route and optimisation case. The **Manage Ports** page is a dedicated three-panel view used to create, edit and curate the port database – a searchable list, an interactive Leaflet map showing every port currently matching the filter, and a controls panel with a coastal-line overlay toggle.

Fig. 6: Port creation.

Reference routes

A reference route defines a corridor across the ocean within which a single optimisation graph will be constructed. Any number of reference routes can exist for the same port pair – one swinging north of a typical storm belt, another hugging a great-circle track, a third following a known traffic-separation scheme. Four creation methods are exposed:

- **Auto Generate:** Backend computes one or more short-corridor options between the two ports asynchronously; live status panel polls every 5 s.
- **Add Manually (map):** Interactive Leaflet editor with click-to-insert / drag / delete waypoints and undo.
- **CSV import:** Comma-separated lat/lon pairs.
- **Import RTZ:** Full RTZ schema parsed; XTD corridors, speed labels and turn-radius circles rendered, and metadata persisted with the route.

Fig. 7: Reference routes between Rotterdam and Lianyungang.

Configuring an optimisation case

Before launching a run, the planner assembles an *optimisation case* – a saved, reusable record that captures every parameter the solver needs. Mandatory fields are the vessel (which determines which power model is loaded), the departure window (ETD – ETA), the origin port and one or more reference routes. A rich set of optional knobs tune the search precisely:

- **Speed range:** Minimum, maximum and step (kt) – discrete speed palette tested on every leg.
- **Departure-time step:** Controls how many ETD offsets are explored within the ETD–ETA window. ETA margin relaxes the deadline by configurable hours.
- **Weather source:** Copernicus historical, ERA5 or ECMWF operational; each component (waves / wind / current) can be toggled independently.
- **Advanced:** Max speed change between consecutive legs, number of Pareto solutions stored per time slot, search fidelity (open sea vs. coastal), custom weather-limit rules.

Fig. 8: Case configuration screen.

Status panel

Once the case is saved and started, the user sees the job's every stage in a live status panel: graph construction, solver progress, post-processing – together with overall completion percentage and per-trade-route status. The job can be stopped at any point.

Fig. 9: Live status panel.

The optimisation worker

The real computation happens inside a per-user Azure Container Instance with 16 CPU cores and 64 GB of RAM. It is provisioned the moment the user logs in and is billed only for the seconds it is actually running, then deleted automatically once the user has been idle for long enough – so cost is directly proportional to actual usage.

Power model – ShipX surrogate

Required power and energy calculations originate in ShipX, SINTEF Ocean's long-running hydrodynamic workbench. To make exhaustive DP feasible, EcoRouter introduces a transitional step: an **interpolated surrogate model** generated from millions of ShipX evaluations spanning combinations of wind, wave and current state, vessel speed and heading. During optimisation, leg power is computed on-the-fly by binary recursive interpolation against this look-up table – orders of magnitude faster than calling ShipX directly, with measured accuracy losses well within operational tolerances.

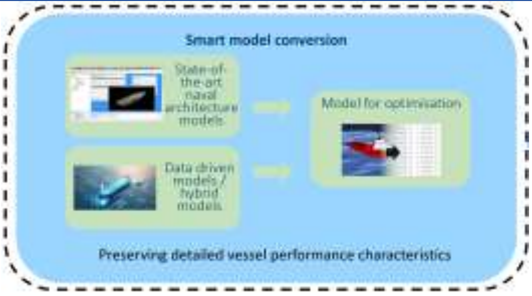

Fig. 10: Vessel model as interpolation file.

Result delivery

Results are persisted the moment the worker finishes. The browser – which has been polling for updates every few seconds – renders them immediately. All Pareto-optimal routes appear on an interactive Leaflet map as coloured polylines, paired with a time-vs-energy scatter chart where each dot is one route option. A collapsible side panel exposes per-route stats and actions (download as RTZ, simulate weather, export raw output to Excel). A *Leg Speeds* toggle annotates each segment with the exact speed assigned by the solver. Weather-data quality is also surfaced on the map, by labelling each graph node with the percentage of missing weather samples.

Fig. 11: Output visualisation: per-route stats (left) and full result view (right) with map, Pareto set and leg-speed labels.

Departure-time sensitivity

In Analysis mode, the same machinery sweeps an entire user-defined range of departure offsets, solves each as a full optimisation and aggregates them into a single chart showing how the optimal energy cost evolves with departure time across days, weeks or a whole year of historical weather.

Fig. 12: Analysis-mode output: Santander → New York, 1-year window, 1-day departure-time step.

Weather playback

For any chosen route, the planner can replay the actual weather field that the vessel encountered. Wave height, wind vectors and ocean currents are rendered as a time-stepped spatial overlay; a vessel marker advances waypoint-by-waypoint with bearing updated to match heading. Heatmap, directional-arrow and numeric-grid layers can each be toggled independently.

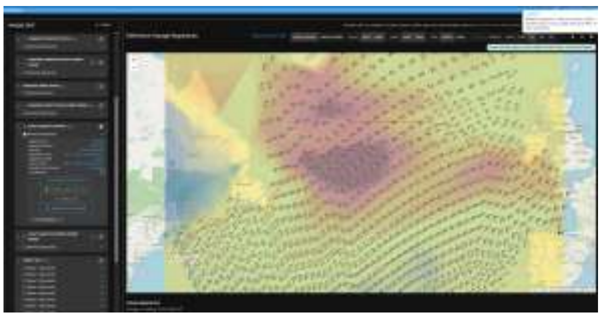

Fig. 13: Dynamic weather visualisation.

Admin: Users & worker fleet

The Manage Users page is an administrator-only operational view that auto-refreshes every 30 seconds. Each row shows online state (green / amber / grey), email, role, last-seen time, the live status of that user's ACI worker container (STARTING / READY) and how long ago it was provisioned. A one-click *Terminate* action stops the container via the Azure SDK and resets the worker DB record, freeing compute resources without waiting for the idle-timeout cycle.

Fig. 14: User and worker-fleet administration.

Admin: World-map land mask

When the solver builds a graph, it tests each candidate arc against a global land mask; any arc that crosses land is removed before the search begins. The World Map editor lets an administrator edit individual coastlines without replacing the whole dataset: drag a vertex to adjust, click an edge to insert a new vertex, right-click to delete; full undo/redo. Drawing entirely new polygons is supported. Edited polygons are exported as a shapefile ZIP, saved to the Azure Files *polygons* share, then promoted to PRODUCTION_WORLD_MAP – the version every worker container loads at startup.

Fig. 15: World-map editor.

Admin: Vessels

Every optimisation case requires a vessel. A vessel record holds two UUIDs that connect the database entry to the external ship-model ecosystem. The first ID points to the **interpolation / SeaWay-format power-model file** that EcoRouter's own interpolation algorithm uses to compute leg power on-the-fly during optimisation (fast, surrogate-driven). The second ID points to the corresponding ShipX *model.zip + propulsion.json*, used for **direct** ShipX **calculation** when high-fidelity results are needed (e.g. weather playback simulation along a finalised route). Keeping both IDs accurate is therefore a hard dependency of both the optimisation workflow and the post-run simulation workflow. The Manage Vessels page exposes all records in a single-card accordion view with inline edit, search, per-user filter and per-vessel access-list management.

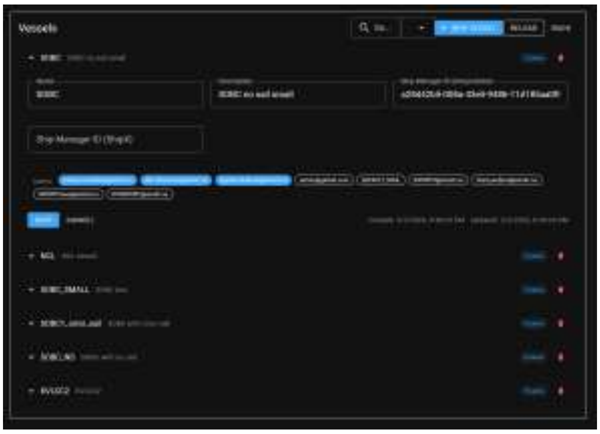

Fig. 16: Vessel management.

Origin and roadmap

EcoRouter started with the EcoRouter IPN project (2022–2024) funded by The Research Council of Norway and industry partners Kristian Gerhard Jebsen Skipsrederi (project owner), Odfjell Tankers and OSM-Thome (previously OSM). It was substantially extended in the Horizon Europe DYNAPORT project, where the optimisation engine was rebuilt and a new web-based GUI was added. The platform builds directly on SINTEF Ocean's long-standing investment in marine hydrodynamics — the same physics that powers ShipX powers every leg-energy calculation in EcoRouter.

Related software

- ShipX – hydrodynamic workbench (vessel power model used as the EcoRouter ground truth).

Contact

Yauheni Kisialiou, PhD
Research Scientist
SINTEF Ocean (Bergen)

yauheni.kisialiou@sintef.no
Phone: +47 90 25 06 10
www.sintef.no/en/software/ecorouter