GNOSIS FREIGHT

WHITE PAPER · GLEC FRAMEWORK

Carbon Emission Tracking

Turning freight emissions into defensible, decision-ready data across ocean, rail, and drayage, calculated from the moves that actually happened, aligned to the GLEC Framework and ISO 14083:2023.

CONTENTS

What this paper covers

01	Why Carbon Reporting Became a Supply Chain Problem	04
02	The Freight Emissions Visibility Gap	06
03	GLEC: A Single Standard for Freight Emissions	08
04	Three Tiers of Carbon Intelligence	10
05	The Missing Ingredient: Operational Data	11
06	The Carbon Intelligence Platform	13
07	From Reporting to Decision-Making	15
08	What Comes Next	16
A	*Appendix · Why Gnosis Is Uniquely Positioned*	17
B	*Appendix · Built to Flex for the Multinational Enterprise*	18

EXECUTIVE SUMMARY

The largest line on the carbon ledger is the hardest one to measure

Upstream transportation is now one of the most consequential and most scrutinized numbers on the corporate carbon ledger. For a typical importer, the freight moving through the supply chain can account for the majority of the company's total emissions, and regulators are no longer treating its disclosure as optional. Yet most organizations still estimate these emissions with annual, spend-based averages and borrowed distances that no auditor can trace back to an actual shipment.

This paper makes the case for a different approach: emissions calculated from the operational data already moving through the Container Lifecycle Management® (CLM) Platform, aligned to the GLEC Framework and ISO 14083:2023, and resolved to the individual container, leg, and node. The result is reporting that holds up under audit today, and becomes an operational input that shapes routing, carrier selection, and network decisions tomorrow.

Maritime transport carries the overwhelming majority of global trade — and the overwhelming majority of its freight emissions.

01

WHY CARBON REPORTING BECAME A SUPPLY CHAIN PROBLEM

The biggest number on the report is the one nobody can measure

For most enterprises, the largest share of the carbon footprint does not come from the buildings they own or the electricity they purchase. It comes from the value chain — Scope 3 — and within Scope 3, upstream transportation and distribution (Category 4) is consistently one of the heaviest contributors. For companies that import goods, particularly those sourcing from Asia, freight emissions frequently rival or exceed the company's entire Scope 1 and Scope 2 footprint combined.

That creates an uncomfortable asymmetry. The emissions that matter most to the disclosure are the ones the company has the least direct visibility into, because the assets producing them (the vessels, the locomotives, the trucks) belong to third parties. The numbers have to be reconstructed from activity data: what moved, how far, by which mode, and with which carrier.

~75%

of corporate emissions sit in the value chain (Scope 3)

Category 4

carries the heaviest freight-reporting load

2026

California SB 253 disclosure begins

Disclosure is no longer voluntary

What was once a reputational exercise is becoming a regulatory requirement. A wave of mandates now reaches across the companies that move goods globally:

FRAMEWORK	WHO IT AFFECTS	REQUIREMENT
EU CSRD	EU-listed companies and large non-EU companies with EU operations; applies to specific carbon-intensive products	Mandatory Scope 3 reporting including Category 4; phased in 2024–2028
California SB 253	Companies with $1B+ revenue doing business in California	Scope 1 and 2 reporting begins 2026; Scope 3 reporting begins 2027
CDP	Investor-requested disclosure across tens of thousands of companies	Logistics emissions required under the Supply Chain module
SBTi	Companies setting science-based climate targets	Scope 3 Category 4 must be included where material to the target
EU CBAM	Importers of certain goods into the EU	Carbon cost and emissions reporting applied at the border

> The common thread: each framework expects defensible, auditable numbers, not estimates a sustainability team assembled by hand the week before a filing deadline.

Why this lands on the supply chain team

Sustainability and ESG teams own the disclosure, but they cannot produce the freight numbers without the supply chain. The activity data that drives Category 4 — origins and destinations, modes, carriers, container weights, actual routings — lives in transportation and logistics systems, not in the sustainability function. In practice, carbon reporting has quietly become a supply chain problem: the quality of the disclosure is now a function of the quality of the operational data the supply chain can produce.

02

THE FREIGHT EMISSIONS VISIBILITY GAP

Spreadsheets, lane averages, and borrowed distance

Most companies reporting freight emissions today do it with some combination of spreadsheets, industry-average emission factors, and assumed routes. A typical workflow looks like this: container data is exported by hand, pasted into a lane-distance worksheet, multiplied by a default factor, and rolled up into an annual figure. It is laborious, slow, and fragile.

The fragility shows up in the details. Default factors get applied to assumed routes, so a shipment is credited with a Panama Canal transit even when the vessel actually sailed around the Cape of Good Hope. Ocean, rail, drayage, and port handling are combined from three or four separate tools, each with its own assumptions, so the totals never quite reconcile. And because the numbers are built on averages rather than actual movements, they cannot be traced back to the shipment that produced them.

COMMON STATE TODAY	RECONCILED TO THE MOVE
Spend-based estimates	• Actual distance
Assumed routes	• One record
Three or four tools	• GLEC-aligned
Untraceable	• Audit-ready

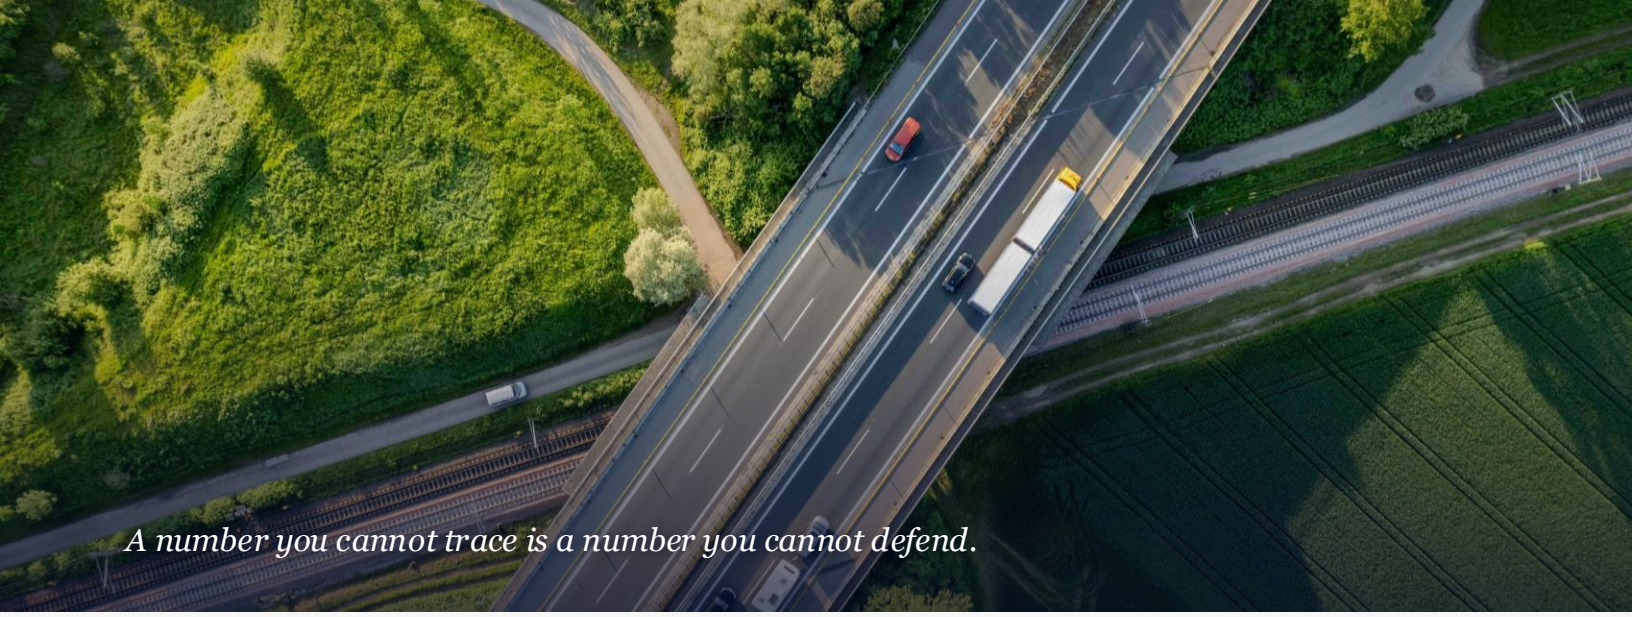

A number you cannot trace is a number you cannot defend.

The audit test

That last point is the one that matters when a disclosure is challenged. An auditor, an investor, or a regulator does not just want a total — they want to know how it was derived and whether it reflects what actually happened. A number that rests on “the industry-average factor for this lane, applied to an assumed route” is difficult to defend. A number that traces to the specific container, the distance it actually traveled, and the mode and carrier that moved it is not.

The cost of the gap

The visibility gap is not only a compliance risk. As long as emissions are an annual, backward-looking estimate, they cannot inform a single operational decision. The company cannot see which lanes are carbon-intensive while there is still time to act, cannot hold a carrier accountable to the sustainability commitments it signed, and cannot factor carbon into the routing and allocation choices it makes every week. The gap keeps emissions in the role of a report to be filed rather than a metric to be managed.

03

GLEC: THE STANDARD THAT MAKES NUMBERS DEFENSIBLE

One method, every mode

The Global Logistics Emissions Council (GLEC) Framework exists to close exactly this gap. It is a single, internationally recognized method for calculating greenhouse gas emissions across every freight mode (ocean, air, road, and rail) plus the handling that happens at ports, terminals, and distribution centers. Its current version, GLEC Framework v3, is fully aligned with ISO 14083:2023, the international standard for quantifying and reporting logistics emissions.

GLEC measures on a well-to-wake (WTW) basis, capturing both the emissions from burning the fuel (tank-to-wake) and the emissions from producing and delivering that fuel (well-to-tank). That gives a complete picture of the climate impact of a move, not just the portion that comes out of the smokestack.

FOUR MODES, ONE ACCOUNTING METHOD

Ocean

The largest share of freight emissions by volume — shaped by vessel, fuel, and load.

Air

The highest emissions per tonne-kilometre — reserved for time-critical freight.

Road

Highly variable — driven by truck size, load, and routing.

Rail

Among the most carbon-efficient land modes for moving containers.

Why GLEC became the standard

GLEC's authority comes from the institutions that have adopted it. It is the methodology recommended by CDP, accepted by the Science Based Targets initiative, compliant with CSRD reporting requirements, and aligned with the GHG Protocol's guidance for Category 4. When a sustainability team asks “is this an accepted standard?”, GLEC is the answer that satisfies every major framework at once — which means a single calculation can serve every disclosure obligation rather than being redone for each.

How GLEC compares to the alternatives

GLEC is among the most complete and most widely accepted. The table below situates it against the methods companies most commonly encounter.

METHOD	HOW IT WORKS	ACCURACY
Spend-based estimation	Apply an industry-average $/kg factor to freight spend (GHG Protocol Tier 1).	Very low — no visibility into mode, distance, or carrier
Carrier-reported	A single carrier reports CO_2e per shipment from its own fleet data.	Medium — accurate for that carrier, not standardized across carriers
Clean Cargo (BSR)	Shipping lines submit trade-lane CO_2 intensity, applied to TEU-km.	Medium — ocean only; a carrier-level average
EcoTransIT World	European modeling platform aligned with EN 16258 / ISO 14083.	Medium-high — requires manual data input
GLEC v3 / ISO 14083	Multimodal standard across ocean, rail, road, and hubs; WTW; primary data preferred.	**High** — per-leg, per-node, audit-ready

The distinguishing feature is scope. GLEC is the only widely accepted methodology that covers the full door-to-door chain (ocean, rail, drayage, and the handling nodes in between) inside one consistent framework. That is what makes it possible to report a complete, comparable number rather than an ocean-only figure with everything else estimated separately.

04

THREE TIERS OF CARBON INTELLIGENCE

From basic compliance to operational precision

GLEC is designed to work with the data a company actually has, and to get sharper as that data improves. The same shipment can be calculated at three levels of precision, each inheriting the one below it and tightening the estimate. The principle is simple: the closer the inputs sit to physical reality, the closer the estimate sits to the truth.

TIER 1 · BASIC

Origin & destination (UN/LOCODE); shipment type

A compliant, preliminary estimate using the shortest feasible route across each mode

TIER 2 · ENHANCED

Vessel / aircraft identity; shipment date; carrier (SCAC); transshipment detail

An estimate that reflects the specific equipment and the specific journey taken

TIER 3 · PRECISION

HIGHEST PRECISION

Container size & type (TEU, dry vs. reefer); cargo weight; fuel type; load characteristics

A calculation that reflects the physical specifics of what actually moved

EACH TIER INHERITS THE ONE BELOW

CLOSER TO GROUND TRUTH

Why the tiers matter

The tiered model means a company never has to wait for perfect data to start reporting. A basic, defensible number is available from day one, and every additional field the supply chain can supply moves the estimate closer to ground truth, without changing methodology or breaking comparability. It also makes the path to improvement legible: a team can see exactly which inputs would sharpen its numbers and prioritize collecting them.

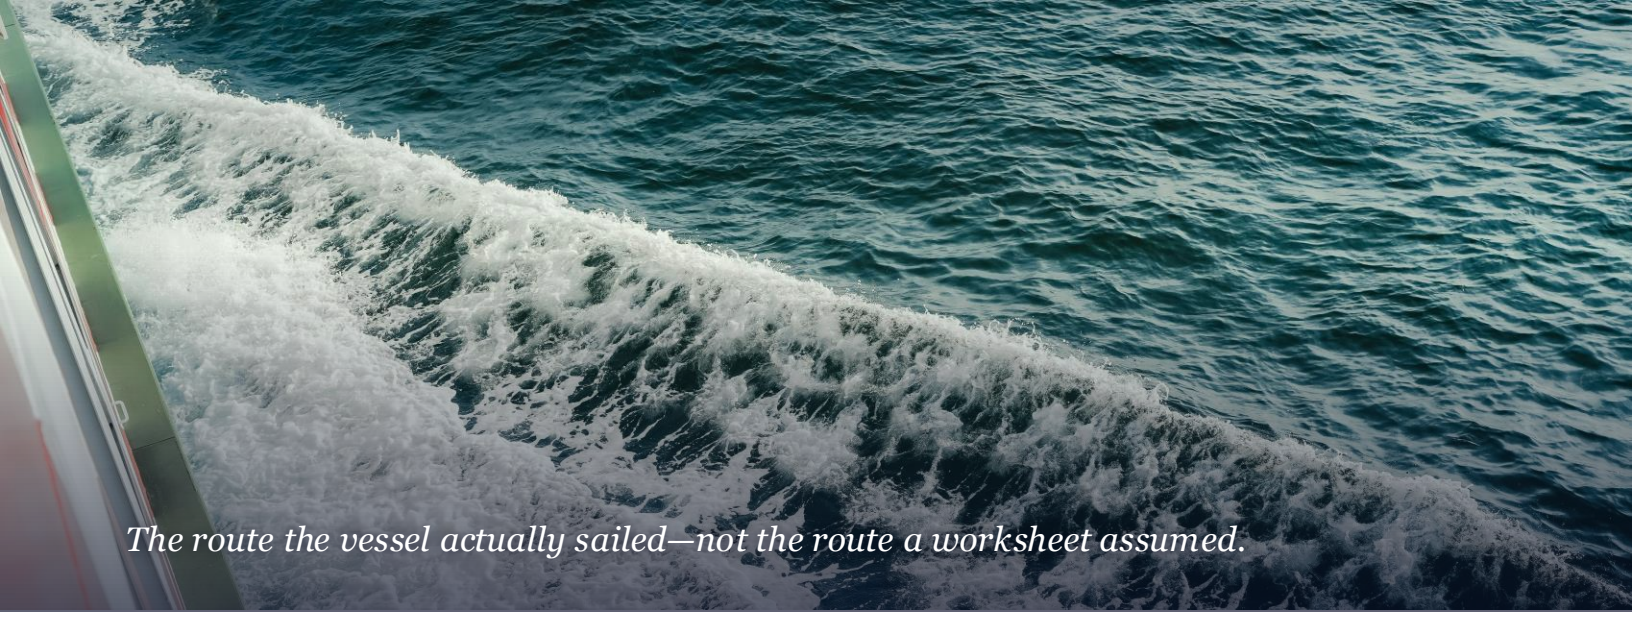

The route the vessel actually sailed—not the route a worksheet assumed.

05

THE MISSING INGREDIENT: OPERATIONAL DATA

Calculated from what actually happened

A standard is only as good as the data fed into it. The reason most freight emissions numbers are weak is not that GLEC is unavailable — it is that companies lack a source of clean, granular, operational data to run through it. This is the gap the Container Lifecycle Management® (CLM) Platform closes. Because CLM already structures and reconciles real-time execution data across global trade, the same record that tracks a container's journey can calculate the carbon that journey produced.

That changes the inputs from assumptions to observations:

Ocean distance	Comes from the actual vessel track, not a theoretical lane: Suez versus Cape, direct versus transshipment, the route the vessel actually sailed.
Rail segments	Are derived from the real rail lines and segments recorded in execution milestones, not a manual mileage lookup.
Drayage	Uses true origin-to-destination road distance from routing data, tied to the dispatched move.
Container weight	Uses the actual per-container cargo weight flowing through CLM, converted to the correct TEU basis.

A disciplined data hierarchy

Following GLEC's data-quality principles, the platform always prefers actual measured data over modeled data, and modeled data over defaults. Each mode has a defined priority order, so every number is built from the best available source, and it is always clear which source was used.

MODE	PRIORITY 1 (BEST)	PRIORITY 2	PRIORITY 3 (FALLBACK)
Ocean	**Actual AIS vessel track**	Previously resolved route for the same lane	Great-circle distance with a documented adjustment factor
Rail	**Rail-network database (actual segments)**	Modeled distance with adjustment factor	—
Drayage	**Road-network routing (actual roads)**	Modeled distance with adjustment factor	—

> Guardrails keep the data honest: when a satellite-derived distance falls outside a plausible range for the expected route, it is rejected and the calculation falls back to the next-best source, so corrupt data cannot inflate a result.

Every leg, every node, every day

The calculation resolves the entire journey inside a single container record — the ocean leg, the intermodal rail leg, the first- and last-mile drayage, and the handling at each port, terminal, and hub — so there is no reconciliation across separate tools. It runs on a daily cadence, so the data is never more than a day old. And it handles in-transit shipments gracefully: a container still at sea receives a projected calculation based on its expected routing, which is replaced by the actual figure once the container is discharged and the real journey is known.

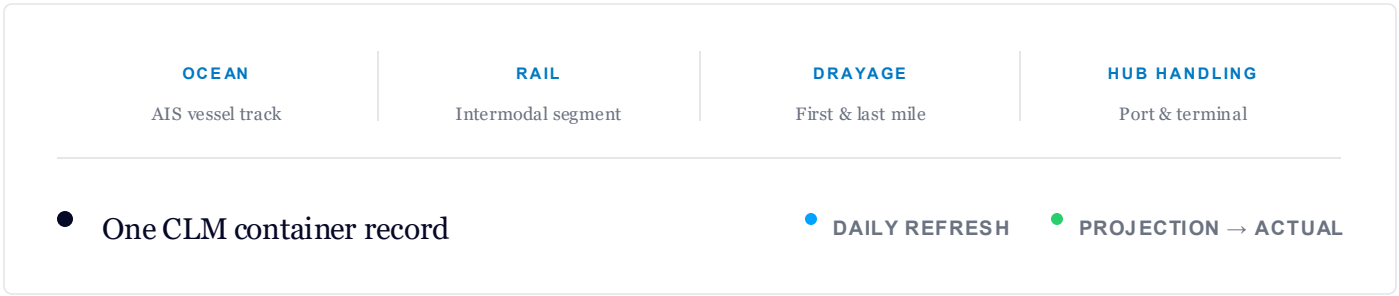

06

THE CARBON INTELLIGENCE PLATFORM

From calculation to working tool

A calculation is only useful if the people who need it can see it, filter it, and export it. The platform surfaces emissions through a set of connected views built for both the sustainability team that has to report the numbers and the supply chain team that produces them.

The carbon dashboard

Figures shown are illustrative, drawn from a sample CLM portfolio.

The dashboard summarizes the network at a glance (total and year-to-date emissions, projected emissions for active containers, average emissions per container, and total shipments) alongside breakdowns of emissions by carrier and by month. Beneath the summary sits a filterable, container-level table that exports straight to a spreadsheet, giving a sustainability analyst the line-item detail needed to assemble a Scope 3 package.

4,726	1,803	497	2.70	6,727
tonnes CO_2	tonnes CO_2	tonnes CO_2	tonnes CO_2	containers
Last Year Emissions	**Current YTD Emissions**	**Projected, Active Containers**	**Average per Container**	**Total Shipments**

Emissions (tonnes) by Ocean Carrier

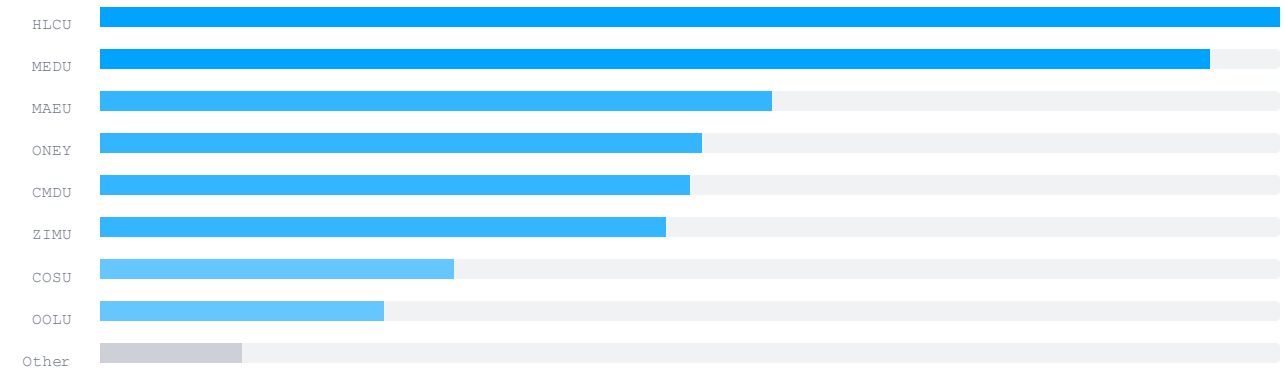

Total Emissions (tonnes) by Month

The network carbon map

The network map renders every lane and every container colored by emissions intensity, so the heaviest-emitting corridors are visible immediately and can be drilled into for container-level detail. Emissions roll up by transport mode and surface the top-emitting legs, turning an abstract total into a map of where the carbon actually is.

18,182

t CO_2 · 6,734 containers · 100% mapped

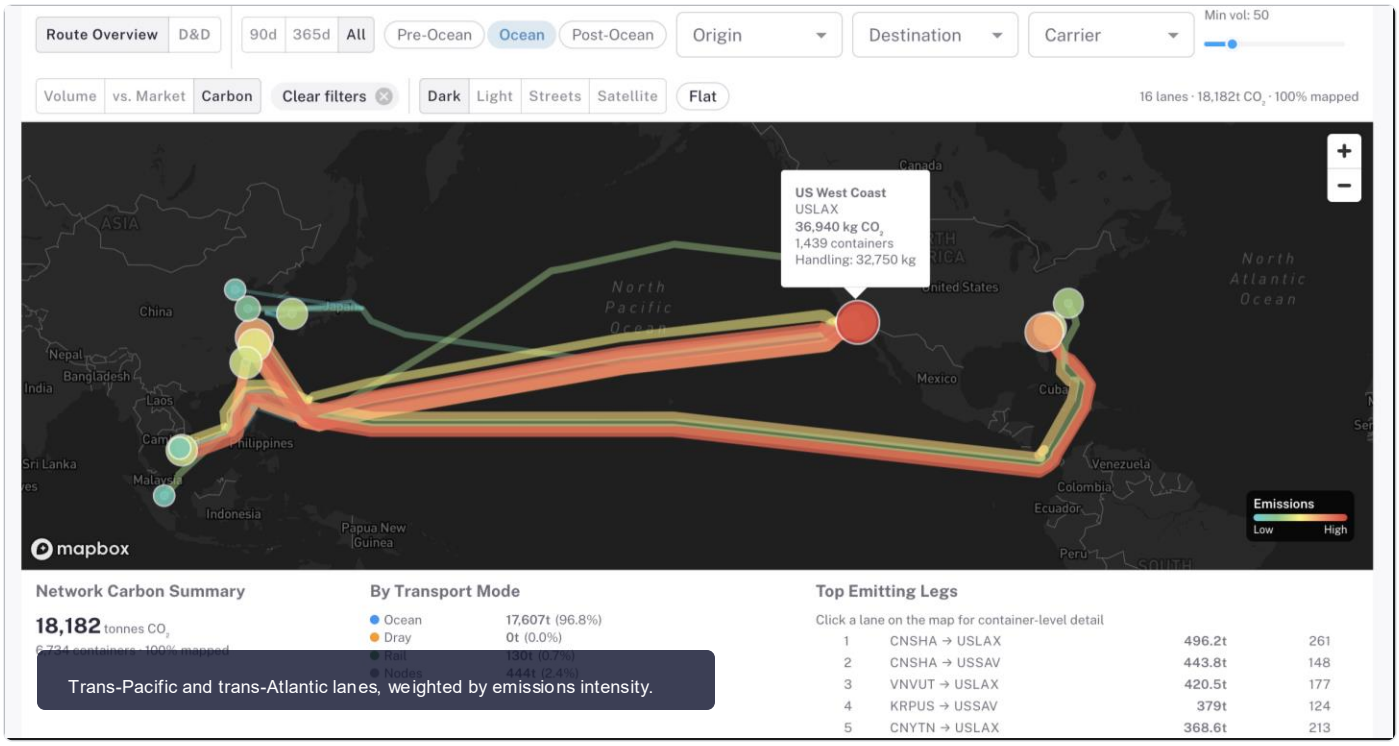

By transport mode

Ocean	17,607t	96.8%
Nodes	444t	2.4%
Rail	130t	0.7%
Dray	0t	0.0%

Top emitting legs

1	CNSHA → USLAX	496.2t
2	CNSHA → USSAV	443.8t
3	VNVUT → USLAX	420.5t
4	KRPUS → USSAV	379.0t
5	CNYTN → USLAX	368.6t

Built for the report and beyond

Every output is structured for audit: GLEC-aligned methodology, traceable to the underlying move, and exportable in the format a consultant or ESG team can drop directly into a disclosure. The same data is available by API for organizations that feed emissions into their own sustainability reporting systems.

07

FROM REPORTING TO DECISION-MAKING

When emissions become an input

Defensible reporting is the first step. Once the numbers are in place and they hold up under audit, the same data can start driving decisions instead of just satisfying disclosures. Emissions stop being a report a company files and start being an input it acts on.

What the data makes possible

With per-container, per-lane, per-carrier emissions refreshed daily, several decisions that were previously impractical become routine:

Lane selection

See which corridors burn more carbon for the same freight, and weigh that into routing.

Carrier scorecards

Measure whether carriers are actually meeting the sustainability commitments they signed, using consistent, comparable numbers rather than self-reports.

Volume allocation

Direct containerized volume toward lower-carbon options when carbon reduction is a stated goal, the same way cost and transit time are already weighed.

Network reduction targets

Set and track emissions-reduction targets at the network level, with the granular data needed to know whether they are being met.

The same data, two jobs

This is the payoff of calculating emissions from operational data rather than annual averages: the figure that defends the disclosure is the same figure that informs the decision. Because the carbon model lives alongside the execution data already used for trade management, emissions can feed the routing and allocation choices a supply chain team makes anyway, without a parallel process, a separate tool, or a six-month data project.

STEP 1

Defensible report

→

STEP 2

Carrier scorecards · lane selection · volume allocation

→

STEP 3

Network-level reduction targets

08

WHAT COMES NEXT

The trajectory

Carbon intelligence follows a natural arc: first make the numbers defensible, then make them operational. The foundation (GLEC-aligned, audit-ready, per-container emissions calculated from real movements) is available now. The next phase deepens the connection between the carbon model and the day-to-day decisions of trade management, so emissions become a standing input to routing, carrier selection, and allocation rather than a quarterly retrospective.

Build it with us

This next phase is being shaped together with the organizations that will use it. There are two ways to take part.

FOR SHIPPERS AND BCOS

You get GLEC-aligned, defensible emissions calculations now, so you are building an audit trail from day one, whether or not your internal team is ready to consume it yet. You also get a direct line into what comes next: the integrations, the lane-level analysis, and the allocation logic. Early participants shape what ships.

FOR CONSULTANTS AND SUSTAINABILITY OFFICERS

Much of the calculation and aggregation work in a Scope 3 program lands on the people who have to defend the numbers. We want those practitioners in the room as the methodology and outputs are refined, so the reports the platform generates match what they actually have to stand behind.

A

APPENDIX

How CLM Data Powers the Calculation

Owned data infrastructure, end to end

Most carbon reporting tools sit on top of data they do not control. They ingest shipment files, route data, and tracking feeds from third-party providers, then apply emission factors on top, which means the quality of the final number depends on systems and assumptions the vendor cannot see into or correct. When an input is wrong, or a provider changes how it models a route, the reporting company inherits the error without a clear way to trace or fix it.

Gnosis is different because the carbon calculation runs on the same data infrastructure that already powers Container Lifecycle Management® Platform. The vessel tracking, the route resolution, the milestone and rail data, and the container records are nested within the platform's native data infrastructure. For an enterprise that must defend its numbers, that ownership has practical consequences:

Control over the inputs	Because we operate the tracking and execution layer, we can see exactly what data fed each calculation and correct it at the source when something looks wrong, rather than waiting on a third party to reconcile it.
Control over the calculation	The methodology, the data-source priority order, and the guardrails are all maintained in one place, so a figure can be traced from the reported total back to the specific move that produced it.
One provider, one record	Ocean, rail, drayage, and hub handling resolve inside the same container record, so there is no combining of separate vendors' outputs, and no reconciliation gaps where their assumptions disagree.

Why this matters for a defensible disclosure

For a large enterprise, the value of owning the infrastructure is auditability. When a regulator or auditor asks how a figure was derived, the answer does not stop at “our vendor provided it.” It traces all the way down to the container, the route, and the data source used — because the same company that calculated the emissions also captured the underlying movement. Control over the inputs and the calculation is what lets the enterprise stand behind the number, not just publish it.

B

APPENDIX

Built to flex for the multinational enterprise

Every company reports carbon differently

No two large enterprises manage carbon the same way. Reporting boundaries, business-unit structures, the frameworks they answer to, the way transportation emissions roll up into aggregated corporate reporting, and the level of granularity their auditors expect all vary, often within a single company across regions and divisions. A rigid, one-size-fits-all output rarely survives contact with a real multinational reporting process. The carbon intelligence platform is built to accommodate that variation rather than fight it.

Tailored outputs	Reports can be shaped to the categories, groupings, and breakdowns a given organization uses (by carrier, lane, business unit, region, or date range) and exported in formats that drop into an existing Scope 3 package or sustainability reporting system, including via API.
Configurable to the enterprise's structure	The data can be aggregated to match how the company actually consumes it, so the figures feed the rollups and disclosures the team already maintains rather than forcing a parallel process.
Aligned to multiple frameworks at once	Because the methodology is GLEC / ISO 14083-based, the same underlying data can support the different frameworks a multinational reports against without recalculating from scratch for each.

The team behind the data

Flexibility depends on people as well as software. Meeting regulatory reporting needs at enterprise scale takes people who understand both the data and the disclosure. Our team works directly with an enterprise's sustainability and ESG functions (and with the external consultants who often carry much of the calculation and aggregation work) to make execution seamless: defining the reporting boundaries, shaping the outputs to what each framework requires, and refining the methodology and reports so they match what the team actually has to defend. The goal is that the data arrives in a form the enterprise can use directly, with the support to put it to work.