

NEPHELE

nephele.earth

Ship-Integrated Sea Salt Aerosol Injection

Restoring Lost Cloud-Albedo Forcing via Exhaust-Plume-Entrained Seawater Spray
Using Existing Scrubber Infrastructure on Commercial Vessels

WHITEPAPER | APRIL 2026 | V6.1

Prepared by Nephele • For Public Distribution

This document describes a commercially viable, regulation-ready method for restoring the cloud-albedo cooling effect lost following IMO 2020 sulfur regulations, using benign saltwater aerosol delivered through existing ship exhaust infrastructure.

I. Summary

For over a century, sulfur dioxide from marine fuel oxidized in the atmosphere to form sulfate aerosol that seeded brighter, more reflective marine clouds. The effect was measurable from space: bright linear cloud features, known as ship tracks, trailing behind vessels across every major ocean corridor.^[27]

In January 2020, the IMO sulfur cap cut shipping sulfur emissions by approximately 80%.^[42] This was an unambiguous public health success. However, it also eliminated the sulfate aerosol that had been providing an unintended global cooling mechanism. The result: an immediate, measurable warming of +0.06 to +0.2 W/m² globally, equivalent to one to three years of additional greenhouse warming arriving all at once.^[5,10,22,43]

Ships did not stop brightening clouds entirely. Even with cleaner fuel and scrubbers, post-2020 vessels still perturb clouds at approximately one-third of pre-2020 levels.^[7] The cloud-brightening capacity was weakened, not destroyed. Sulfur in emissions continue to act as cloud condensation nuclei, or CCN; we must simply replace the lost radiative forcing with a cleaner particle.

Nephele aims to supplement the remaining two-thirds gap with pure, atomized seawater. A bolt-on module installed above the exhaust scrubber (already in operation on the majority of new cargo ships) uses a patent-pending rotary splash atomizer to produce fine aerosol from coarser nozzle spray, with a mist eliminator passing only sub-3 μm droplets, those in the CCN-active size range, into the warm, water-saturated exhaust plume.

What happens next is a core insight of this proposal. The injected droplets do not dry. Post-scrubber exhaust is warm and saturated with water vapor. Observational datasets from marine stratocumulus regions show that boundary layer relative humidity exceeds the NaCl deliquescence point of 75% for roughly 85 to 95% of the time at all altitudes below cloud base.^[3,4] The seawater droplets remain as concentrated salt solution droplets throughout their brief and direct transit to cloud base, concentrating during plume dilution but not crystallizing under the humidity conditions that prevail in target corridors. We do not rely on passive convective lofting – warm ship plumes exist as a free and abundant transport mechanism. No drying step, no deliquescence barrier, no activation delay. The droplets that reach cloud base are effectively pre-activated cloud condensation nuclei, remaining activated at very low supersaturations typical of marine stratocumulus. This method piggybacks entirely on existing scrubber and exhaust architecture, requiring minimal engineering integration.

This bypasses the central challenge of every other marine cloud brightening proposal: the spray-dry-loft-deliqesce-activate chain.^[24,30,40] Each step imposes losses. Nephele eliminates all of them, while using infrastructure that already exists.

Over 4,000 scrubber-equipped vessels are in operation today. Nephele will restore ship tracks across the global ocean by supplementing weakened post-2020 radiative forcing with pure seawater aerosol delivered through the ship's own exhaust plume, under IMO MEPC oversight, in designated open-ocean corridors, with satellite and shipboard verification.

2. The Gap

Prior to 2020, global shipping emitted roughly 10 Tg SO₂ per year, approximately 13% of all anthropogenic sulfur dioxide emissions.^[42] This SO₂ oxidized in the marine boundary layer, primarily via aqueous-phase reaction with H₂O₂ inside cloud droplets, to form sulfate aerosol that enhanced CCN concentrations by roughly an order of magnitude above the clean marine background, from approximately 100 to 120 cm⁻³ to 500 to 1000+ cm⁻³.^[2,33]

The 2020 sulfur cap reduced these emissions by approximately 80%. Diamond and Boss (2025) used shipping diversions around the Cape of Good Hope, where Houthi attacks on Red Sea shipping created a “natural experiment” in traffic density, to quantify the impact: the sulfur cap produced a ~67% reduction in cloud droplet number enhancement per unit of shipping.^[7] Post-2020 ships retain approximately one-third of their pre-2020 cloud perturbation capacity. Ship tracks still occasionally form; “invisible” cloud perturbations, too subtle for satellite detection as linear features but statistically measurable when ship positions are known, remain substantial.^[27]

The measured radiative forcing from this reduction is +0.06 to +0.2 W/m² globally.^[5,10,22,43] Hansen et al. (2026) argue that the reduction of shipping aerosols is one component of a broader decline in aerosol cooling that, combined with a climate sensitivity of 4-5°C for doubled CO₂, is driving the extraordinary acceleration of sea surface temperature warming observed since 2023.^[12] Detectable ship tracks account for only about 5% of the total aerosol indirect forcing from ship emissions; the majority comes from diffuse, corridor-wide perturbations.^[44]

The engineering challenge is supplementation, not full replacement. Post-2020 ship exhaust plume itself already provide a strong CCN foundation. Nephele closes the remaining two-thirds gap, approximately 270 to 600 additional CCN cm⁻³ within the plume cross-section, with pure seawater.

3. Plume Delivery Advantage

Every marine cloud brightening concept proposed to date follows the same pathway: spray seawater at sea level, evaporate the droplets to dry salt crystals, and rely on marine boundary layer turbulence to loft those crystals to cloud base over 30 to 60 minutes.^[24,30,40] Every step leaks particles.

Evaporative cooling at the spray source creates negative buoyancy that suppresses vertical transport.^[21] Dry crystals must survive gravitational settling, coagulation, and dry deposition during the long passive transit. At cloud base, dry NaCl crystals must first deliquesce, absorbing water at ~75% RH, before they can activate as CCN. Deliquescence requires energy, thus not all dry particles delivered to cloud base using these methods become useful.

Nephele's architecture sidesteps this entire chain.

3.1 Plume Transit: Stack to Cloud Base

At the stack (0 m). Post-scrubber exhaust exits at roughly 50 to 80°C and 100% relative humidity. Water vapor content is approximately 150 g/kg, compared to ambient air at about 10 g/kg. The injected seawater droplets enter this warm, saturated environment. They cannot evaporate because the surrounding air is already saturated. The plume is buoyant because it is warm relative to the ambient marine boundary layer.

First 50 to 100 m. The plume rises rapidly. Ambient air at roughly 18°C and 80% RH is entrained at the plume edges, but the plume core remains close to saturation because the moisture excess is so large (a 15:1 ratio of plume to ambient water vapor content). The plume is visible as a dense white condensation cloud because temperature mixing at the boundaries creates local supersaturation, the same physics that makes breath visible on a cold day. Droplets remain at roughly their original size. Plume diameter expands from roughly 3 m to 20 to 30 m.

100 to 300 m. Entrainment has diluted the plume significantly. The unmixed fraction is roughly 10 to 30%. Two competing effects govern RH: the plume is getting drier from entrainment of 80% RH air, but it is also cooling with altitude at the moist adiabatic lapse rate (~6.5°C/km), which pushes RH back up. Net result: RH in the plume core remains above 85 to 90%, with 80 to 85% at the edges. Droplets at the plume edges begin to concentrate. A 1 µm seawater droplet at 85% RH equilibrates to about 0.4 µm. It has lost water, but its dissolved NaCl mass is unchanged. It is now a smaller, more concentrated droplet with lower settling velocity and higher solute-to-volume ratio.

300 to 700 m (approaching cloud base). The plume is thoroughly mixed with ambient air. RH is now set by the ambient MBL humidity profile, which in stratocumulus regions rises from about 80% at the surface toward saturation (100%) at cloud base. So as the plume dilutes and loses its own moisture identity, the ambient RH is *increasing* with altitude. The droplets that concentrated during mid-transit now slowly re-absorb water as RH climbs toward cloud base.

At cloud base. Supersaturation develops in rising air parcels. The droplets, which are concentrated NaCl solution rather than dry crystals, activate at very low supersaturations typical of marine stratocumulus. There is no solid-to-liquid transition, no deliquescence kinetics, and no wetting delay. They are already liquid and already carrying sufficient solute mass to serve as CCN without the kinetic barriers that dry crystals must overcome.

The full transit follows a V-shaped trajectory in concentration: most dilute at the stack, most concentrated during mid-transit as the plume mixes with drier ambient air, then diluting again as ambient RH rises toward cloud base. In stratocumulus regions, ambient MBL humidity exceeds the NaCl deliquescence point of 75.3% roughly 85 to 95% of the time.^[3,4] Real plumes are turbulent and intermittent, with shear-driven filaments and transient mixing at the edges, and brief excursions below 75% RH cannot be categorically excluded. However, the mean-state humidity profile in the target regions strongly favors solution-phase persistence throughout transit. The mid-transit concentration is actually beneficial: smaller droplets have lower settling velocity and reduced coagulation probability, making them easier to transport and less likely to merge with neighbors during the remaining ascent.

3.2 Pre-Activated CCN

In ambient MCB proposals, dry salt crystals must overcome two barriers at cloud base: deliquescence (the solid-to-solution transition at ~75% RH) and then growth past the critical wet diameter governed by Köhler theory.[31] Both steps impose kinetic delays and require sufficient ambient supersaturation.

Nephele's droplets face neither barrier. They arrive at cloud base as liquid solution, already past deliquescence, already larger than their critical wet diameter. There is no activation step. They are CCN the moment they enter the cloud.

The only scenario where activation matters is the edge case: a droplet that dries during a transient sub-75% RH excursion mid-transit (Section 3.5). Even then, the solute mass is more than sufficient for reactivation. A 0.5 μm wet seawater droplet carries enough NaCl to reactivate at $Sc = 0.028\%$; a 1.0 μm droplet at $Sc = 0.010\%$. Marine stratocumulus cloud base supersaturation is 0.1 to 0.3% — roughly 10 to 30 times higher than required. Reactivation is not a limiting factor.

3.3 Thermal Lofting

The warm exhaust plume provides buoyant vertical transport. Initial vertical velocities at the stack are on the order of 5 to 15 m/s, decaying as entrainment reduces the temperature differential. Estimated transit time to cloud base (500 to 1000 m) is on the order of 5 to 15 minutes, depending on initial buoyancy flux, ambient stratification, and wind shear.^[14] This is substantially faster than the 30 to 60 minutes required for passive MBL turbulent lofting in ambient MCB proposals, though the exact advantage depends on conditions that will be measured directly during the Phase 0 campaign. The settling velocity of a 3 μm seawater droplet is approximately 0.3 mm/s, negligible compared to plume updraft velocities. Even droplets that have concentrated to 0.35 μm during mid-transit settle at roughly 4 $\mu\text{m/s}$, which is functionally zero compared to any vertical air motion in the MBL.

3.4 Comparison to Ambient MCB Injection

Plume delivery puts a substantially higher fraction of injected particles at cloud base as effective CCN. Ambient MCB injection must contend with evaporative cooling (which suppresses vertical transport at the source^[21]), passive lofting losses over 30 to 60 minutes, complete drying to solid crystals, and deliquescence kinetics at cloud base. Conservative estimates of end-to-end delivery efficiency for ambient injection range from 0.5 to 10.6%^[14]. Based on the physical arguments in Sections 3.1 through 3.3, we hypothesize that Nephele's plume delivery achieves 70 to 95% delivery efficiency. This is the central testable claim of the proposal and the primary target of Phase 0 measurements. Even at the lower bound, the improvement over ambient injection is roughly an order of magnitude. This delivery advantage is what makes the nozzle requirements tractable.

3.5 Sources of Uncertainty (to be resolved in Phase 0 testing)

The plume delivery advantage is the central physical claim of this proposal. It rests on assumptions that are individually well-grounded but have not been directly measured in a ship exhaust plume carrying injected sea salt. The following uncertainties will be resolved by Phase 0 and Phase 1 measurements:

RH variability in the plume. The argument that droplets remain as solution depends on MBL humidity exceeding the NaCl deliquescence point (~75% RH) for the large majority of plume residence time. Observational datasets support this in stratocumulus regions^[3,4], but real plumes

experience turbulent intermittency, dry intrusion events, and diurnal decoupling that may create transient sub-75% excursions. Phase 0 will measure RH profiles within and around operating scrubber plumes directly.

Plume coherence and rise rate. Transit time depends on initial buoyancy flux, entrainment rate, ambient wind shear, and boundary layer stratification. The estimated 5- to 15-minute range is based on standard buoyant plume models^[14], but real ship plumes are subject to shear breakup, filamentation, and intermittent structure that smooth profiles do not capture. Phase 0 instrumentation (ship-mounted lidar, drone profiling) will characterize actual plume geometry and rise behavior.

Delivery efficiency. The hypothesized 70 to 95% delivery efficiency combines minimal drying losses, no deliquescence barrier, near-complete activation, and faster transport. Each factor is physically motivated but the compound product has not been measured. Phase 0 will measure CCN concentration at multiple altitudes within the plume to directly constrain this number.

Cloud response nonlinearity. The Twomey effect is logarithmic in CCN concentration.^[36] Adding CCN to an already-perturbed plume produces less brightening per particle than adding them to a clean cloud. Post-2020 ships still emit CCN at one-third of pre-2020 levels, so Nephele is adding particles to an already-perturbed plume. Phase 1 satellite analysis will quantify actual albedo response per unit of injected aerosol. Cloud albedo sensitivity to additional CCN decreases as background concentration increases – meaning that ambient conditions and local weather must be accounted as well.

4. Prior Approaches

The marine cloud brightening field has spent decades searching for the perfect spray nozzle. This search has produced extraordinary science and zero deployed systems.

Salter's original design called for *billions* of micro-fabricated silicon nozzles.^[30] Cambridge's Centre for Climate Repair is exploring superheated converging-diverging nozzles with shock-wave atomization. The University of Washington's effervescent nozzle work achieved a count median diameter of approximately 50 nm dry, requiring ~512 W per nozzle at 71% CCN-active fraction.^[13] Scaling to the 10^{15} CCN/s Wood criterion^[40] requires ~814 nozzles drawing ~417 kW, feasible on a research vessel but logistically challenging for unattended commercial operation.

All of these approaches share a common assumption: the nozzle must produce the right particle at the point of emission, because the particle then enters ambient air and must survive the full spray-dry-loft-activate chain. This assumption demands extraordinary nozzle precision. Nephele's plume delivery architecture removes the assumption. Because injected droplets stay as solution droplets and arrive effectively pre-activated at cloud base, the nozzle does not need to produce a narrow sub-100 nm dry distribution. It needs to produce droplets in the 0.3 to 3 μm wet range, a far more achievable target for industrial spray hardware, readily retrofitted to existing exhaust scrubbers.

5. The Nephele Module

5.1 Architecture

The Nephele module is a self-contained bolt-on unit that mounts above the scrubber exhaust stack, above the scrubber's own mist eliminator. It does not modify with the scrubber or affect its MARPOL Annex VI compliance.

The module contains three components: air-assist spray nozzles feeding a patent-pending rotary splash atomizer, a removable mist eliminator above it, and instrumentation/control electronics. Seawater is tapped from the scrubber seawater circuit. No additional seawater intake or filtration is required, as this is already in place for the existing exhaust scrubber.

5.2 The Mist Eliminator

A mist eliminator (or demister) is a mesh of thin wires or vanes with 97 to 99% void fraction.^[1] The capture mechanism is inertial impaction: gas weaves around wires; large droplets cannot follow the turns and impact. Small droplets follow the streamlines and pass through. Sub-micron particles pass with near-100% efficiency. Particles above $\sim 3 \mu\text{m}$ are captured with high efficiency.^[34] Mist eliminators are standard, mature, off-the-shelf technology on every scrubber-equipped ship, as well as being used widely in the oil & gas industry and industrial chemical production.

Pressure drop across a mist eliminator is less than 250 Pa ^[9]. For context, a marine scrubber already imposes several thousand Pa of backpressure on the ship engine exhaust. In an external configuration, where the module sits alongside the stack and the rotor's own fan action drives flow through the eliminator, the module imposes no backpressure on the exhaust system at all. In either configuration, the eliminator is a replaceable cassette swapped during routine port calls.

5.3 Rotary Splash Atomization

Nephele's injection architecture is designed to accommodate multiple atomization technologies as they mature. The atomizer module is the platform; the nozzle type feeding it is a replaceable commodity. The platform (mounting, power, control, verification instrumentation, data handling, regulatory positioning, fleet contracts) is the asset.

Near-term deployment uses commercially available air-assist nozzles feeding a rotary splash atomizer. Air-assist nozzles are catalog industrial components with decades of marine service history. They produce a broad polydisperse spray with a volume median diameter ($Dv50$) of approximately $15 \mu\text{m}$ and geometric standard deviation σ_g of approximately 1.8.^[25] $Dv50$ is the diameter below which 50% of the spray mass is contained; because mass is concentrated in the larger droplets, the count median diameter (CMD) — the diameter below which 50% of droplets by number fall — is much smaller: $\text{CMD} = Dv50 \cdot \exp(-3 \ln^2 \sigma_g) \approx 5.3 \mu\text{m}$. Without secondary atomization, only a small fraction of this output is useful for CCN production (Section 6.4). The rotary splash atomizer is the critical secondary stage that pushes the spray distribution into the useful size range.

Splash-impact atomization is well established: when a spray droplet impacts a wetted surface at high Weber number, the spreading film becomes unstable via Rayleigh-Taylor instability, forming crown structures with fingers that pinch off as secondary daughter droplets.^[6,41] The daughter droplet diameter scales approximately as $We^{-1/4}$ relative to the parent, where We is the impact

Weber number.[38] However, splash efficiency depends critically on the dimensionless film thickness h/D on the impact surface, where h is the liquid film thickness and D is the impacting droplet diameter. Published data establish that efficient splash (full crown formation, fine daughter production) requires $h/D < 1$; at $h/D > 3$, splash is largely suppressed and impacting droplets are absorbed into the film.[6,37]

The film thickness problem. On a static plate receiving continuous spray, the equilibrium film thickness is governed by the Nusselt thin-film equation: $h_{eq} = (3 \cdot \nu \cdot J / a)^{1/2}$, where ν is kinematic viscosity, J is the volumetric spray deposition flux, and a is the drainage acceleration. Under gravity alone ($a = g = 9.81 \text{ m/s}^2$), at a representative deposition flux of $J \approx 0.10 \text{ m/s}$, $h_{eq} \approx 180 \text{ }\mu\text{m}$ for seawater. Even for the largest droplets in the spray distribution, $h/D \gg 3$. For the count median ($5.3 \text{ }\mu\text{m}$), $h/D \approx 34$. Splash is deeply suppressed. This is a fundamental limitation: no static plate geometry, material, or surface treatment can achieve $h/D < 1$ at these throughput rates under gravitational drainage.

Centrifugal film equilibrium (patent pending). Nephele's rotary splash atomizer solves the film thickness problem by substituting centrifugal acceleration for gravity. Air-assist nozzles spray onto the blade faces of a rotor spinning at $\sim 6,000 \text{ RPM}$. At a blade radius of 0.10 m , centrifugal acceleration is $\omega^2 r \approx 39,500 \text{ m/s}^2$, roughly 4,000 times gravitational acceleration. The equilibrium film thickness drops from $180 \text{ }\mu\text{m}$ to approximately $2.8 \text{ }\mu\text{m}$, yielding $h/D \approx 0.5$ for count-median droplets and $h/D \approx 0.2$ for $Dv50$ -sized droplets: well within the efficient splash regime across the full drop size distribution. The blade tangential velocity (63 m/s at mid-span) provides the impact Weber number for splash atomization. This is a fundamentally different device from a conventional spinning disc atomizer, which relies on centrifugal edge breakup to form droplets. In Nephele's rotary atomizer, centrifugal force prevents film accumulation while the blade surface serves as the splash-impact target.

Because blade tangential velocity varies with radius ($v = \omega \cdot r$), the Weber number and daughter droplet size vary across the blade span. For $Dv50$ -sized parents ($15 \text{ }\mu\text{m}$), We ranges from ~ 200 at the hub ($r = 5 \text{ cm}$, 31 m/s) to $\sim 1,900$ at the tip ($r = 15 \text{ cm}$, 94 m/s), producing daughters of approximately $4 \text{ }\mu\text{m}$ at the hub down to approximately $2.3 \text{ }\mu\text{m}$ at the tip. At mid-span ($r = 10 \text{ cm}$, 63 m/s , $We \approx 850$), daughter diameter is approximately $2.8 \text{ }\mu\text{m}$ for $Dv50$ -sized parents. Smaller parents — which dominate the distribution by number — produce finer daughters, and nozzles are targeted at the outer blade span where We is highest. The net result is a central daughter CMD estimate of approximately $2 \text{ }\mu\text{m}$ (detailed in Section 6.4), at which approximately 72% of daughter droplets by number pass the mist eliminator (compared to 16.5% of the original parent spray). Each parent impact produces on the order of 10 to 20 daughter fragments.[6,41] The combined effect of count multiplication and size shift is approximately 30× more useful CCN output than bare nozzles alone.

Additionally, because the nozzles are distributed over a fraction of the rotor circumference, each blade alternates between spray deposition and centrifugal clearing during each revolution. During clearing intervals, centrifugal drainage actively thins the film below the steady-state equilibrium, meaning the Nusselt equation represents a conservative worst case. The rotor's fan action also drives aerosol through the downstream mist eliminator without requiring an external pressure source.

Key operational concerns include salt crystal buildup on the blade surfaces and corrosion from continuous seawater contact. Both are managed by the replaceable cassette design and corrosion-resistant blade materials (Inconel 625 or ceramic). Blade surface roughness and coatings are design variables to be optimized during Phase 0 hardware development.

As effervescent, convergent-divergent, or other precision nozzle technologies mature for continuous marine operation, they plug directly into the same rotary atomizer module as the input spray source. The platform is nozzle-agnostic by design.

5.4 Exhaust Plume does not affect Droplet Size

Secondary breakup occurs when aerodynamic Weber number exceeds $We_{crit} \approx 12$.^[11] For post-scrubber exhaust ($\rho_{gas} \approx 0.9 \text{ kg/m}^3$, $v_{rel} \approx 15 \text{ to } 30 \text{ m/s}$, $\sigma_{water} \approx 0.072 \text{ N/m}$), a $50 \text{ }\mu\text{m}$ droplet gives $We \approx 0.14 \text{ to } 0.56$, far below breakup threshold. What the nozzle and rotary atomizer produce and the mist eliminator passes is what enters the boundary layer.

5.5 Installation, Maintenance, and Cost

Scrubber systems are subject to regular MARPOL Annex VI inspection. The scrubber stack exit is routinely accessed for maintenance. Installation of the Nephele module can be performed during a scheduled port call. Nozzle replacement is a hand-tool operation. Mist eliminator cassette swap takes less than an hour.

Estimated per-vessel installation cost: approximately \$250,000 including hardware, installation, and commissioning. For context, a scrubber installation costs \$2 to 8 million. Dedicated cloud brightening proposals require purpose-built vessels — Nephele uses 4,000+ ships that already exist.

6. Particle Physics and CCN Production

6.1 Wet Droplet to Dry Crystal: Size Relationship

When a seawater droplet fully evaporates, the dry NaCl crystal is smaller than the original droplet by a factor determined by salinity and density. For seawater at 3.5% salinity ($\rho_{sw} = 1025 \text{ kg/m}^3$, $\rho_{NaCl} = 2160 \text{ kg/m}^3$): $d_{dry}/d_{wet} = (\rho_{sw} \times w_s / \rho_{NaCl})^{2/3} \approx 0.255$. A 3 μm droplet dries to $\sim 765 \text{ nm}$. A 1 μm droplet dries to $\sim 255 \text{ nm}$. A 0.5 μm droplet dries to $\sim 128 \text{ nm}$.

However, in Nephele's architecture, the droplets do not fully dry. As described in Section 3, the marine boundary layer RH remains above the NaCl deliquescence point for the large majority of transit time. The droplets equilibrate to concentrated solution droplets whose size depends on ambient RH: at 80% RH, a 1 μm seawater droplet equilibrates to $\sim 0.35 \mu\text{m}$; at 90% RH, to $\sim 0.5 \mu\text{m}$. At cloud base ($\sim 99\%$ RH), the droplets are near their original injected size. The dry crystal diameter is relevant only as a measure of the solute mass each droplet carries, and all droplets above $\sim 0.3 \mu\text{m}$ wet carry more than enough solute for CCN activation.^[31]

6.2 NaCl vs. Sulfate as CCN

Sodium chloride is a more effective CCN than the sulfate aerosol it replaces. NaCl has a hygroscopicity parameter (κ) of 1.28, compared to 0.61 for ammonium sulfate.^[31] Particle for particle, NaCl activates at lower supersaturation — or equivalently, a smaller NaCl particle does the same job as a larger sulfate particle. And unlike sulfate, which had to nucleate from SO_2 gas, survive atmospheric transport as a dry crystal, and deliquesce before activating, Nephele's NaCl arrives at cloud base already in solution. Better material, better delivery.

Additionally, the input material is seawater, something that belongs in the ocean, is chemically inert in the atmosphere, and ecologically benign. The sulfate system it replaces was a chaotic, uncontrolled byproduct of burning dirty fuel that acidified any environment it touched and damaged ecosystems worldwide. Nephele replaces a damaging pollutant with the ocean itself.

6.3 Controllability: The Strategic Advantage

Pre-2020 sulfate cloud brightening was an entirely uncontrolled side effect of burning dirty fuel. SO_2 left the stack as a reactive gas and entered a complex atmospheric chemistry chain: gas-phase oxidation by OH radicals (slow, daytime only), aqueous-phase oxidation by H_2O_2 and O_3 inside cloud droplets (fast but intermittent), nucleation of new particles or condensation onto existing ones, and eventual activation as CCN. Nobody controlled the particle size distribution, where the conversion happened, whether receptive clouds were overhead, or whether the resulting aerosol would brighten or darken the clouds. The sulfate system was a brute-force atmospheric reactor running unsupervised across the entire ocean.

Sea salt aerosol does none of this. NaCl is chemically inert in the marine boundary layer. It does not oxidize, does not nucleate new particles through gas-phase chemistry, and does not react with atmospheric oxidants to form secondary aerosol. A NaCl solution droplet injected at the stack is the same NaCl solution droplet that arrives at cloud base. Its size, solute mass, and CCN activation properties are known at the point of injection and do not change during transit.

This means Nephele's forcing is predictable, tunable, and verifiable in a way that sulfate forcing never was. We can turn injection on or off in response to real-time satellite cloud data. We can

modulate injection rate based on observed meteorological conditions. We can deactivate in regions where cloud response is unfavorable. The pre-2020 sulfate system had none of these capabilities. Our per-particle forcing may be somewhat smaller than the theoretical maximum, because we are not exploiting the atmosphere as a free particle factory. But we know exactly what we are putting into the air, exactly where, and exactly when. That tradeoff, less raw power but far greater precision and accountability, is what makes ship track restoration a governable climate intervention rather than an uncontrolled accident of fossil fuel combustion.

6.4 Per-Ship CCN Budget

The central question is: how does Nephele's CCN output compare to what a pre-2020 ship was delivering to cloud base? Both sides of this comparison involve multiple uncertain parameters. We walk through each step explicitly below.

Pre-2020 sulfate system: fundamentally different and harder to quantify. A direct particle-for-particle comparison between Nephele and the pre-2020 sulfate system is misleading, because the two systems produce CCN through entirely different mechanisms. Nephele injects finished solution droplets. The pre-2020 system emitted SO₂ gas, which the atmosphere then converted to sulfate particles through complex, uncontrolled chemistry over hours to days.

A pre-2020 vessel burning HFO at 2.7% sulfur emitted approximately 25 kg/hr SO₂. Within the ship track timescale of several hours, roughly 5 to 15% of this SO₂ converted to sulfate aerosol. If all of this sulfate nucleated new 150 nm particles (ammonium sulfate density 1770 kg/m³), the equivalent particle production rate would be approximately 3.3×10^{14} particles/s. But this is an upper bound, not a realistic estimate of new particle formation.

SO₂ conversion to sulfate occurs via two pathways. Gas-phase oxidation by OH radicals (approximately 30% of conversion) can nucleate new particles, but in the polluted marine boundary layer, most of this sulfate condenses onto pre-existing particles rather than forming new ones. Aqueous-phase oxidation by H₂O₂ and O₃ inside cloud droplets (approximately 70% of conversion) adds sulfate mass to existing cloud droplets but does not create new CCN. The fraction of converted sulfate that actually nucleated new particles is poorly constrained and varied along every ship track.

Additionally, combustion produced roughly 1.5 kg/hr of primary particulate matter (black carbon, organic carbon, primary sulfate, ash) with CMD approximately 100 nm. These particles did not depend on atmospheric conversion, but their delivery to cloud base and activation efficiency are themselves uncertain (estimated 10 to 20% delivery, 50 to 70% activation).

Observational ground truth. Rather than attempting to close a bottom-up particle budget with poorly constrained parameters, we anchor the comparison to direct measurements. Diamond and Boss (2025) measured the actual cloud response to the IMO 2020 sulfur cap: a ~67% reduction in cloud droplet number enhancement per unit of shipping.^[7] This measurement integrates all pathways: primary PM, gas-phase nucleation, condensational growth, aqueous-phase modification, and cloud processing. It is the empirically measured deficit that Nephele aims to fill.

Traditional hardware. A single air-assist nozzle operating at 1 L/min (1.67×10^{-5} m³/s) can produce a polydisperse spray with a volume median diameter of approximately 15 μm and geometric standard deviation σ_g of approximately 1.8.^[25] For a log-normal distribution, the count median diameter (CMD) = $D_{v50} \times \exp(-3 \ln^2 \sigma_g) = 15 \times \exp(-3 \times 0.587^2) = 15 \times 0.356 \approx 5.3$ μm. The mean droplet volume is $(\pi/6) \times \text{CMD}^3 \times \exp(4.5 \ln^2 \sigma_g) = 3.73 \times 10^{-16}$ m³. Total droplet production: flow rate / mean volume $\approx 4.5 \times 10^{10}$ droplets/s per nozzle.

The mist eliminator passes droplets below approximately 3 μm wet. For a log-normal distribution with $\text{CMD} = 5.3 \mu\text{m}$ and $\sigma_g = 1.8$, the cumulative fraction below 3 μm is approximately 16.5% by number. Without secondary atomization, each nozzle delivers roughly 7.4×10^9 droplets/s past the eliminator. At 100 nozzles with 85% plume delivery: approximately 6×10^{11} effective CCN/s per vessel. This is orders of magnitude below any estimate of pre-2020 per-ship CCN. Air-assist nozzles alone, without secondary atomization, are entirely insufficient.

Traditional hardware with rotary splash atomization. At 50% splash mass fraction and approximately 15 daughter fragments per parent impact^[6,41], total particle count increases by roughly 8 \times . The daughter distribution is centered near 2 μm wet CMD (from $\text{We}^{-1/4}$ scaling^[38]), so approximately 72% of daughter droplets pass the mist eliminator, compared to 16.5% of the original parent spray. The combined effect of count multiplication and size shift yields approximately 30 \times more useful CCN than bare nozzles. With rotary splash atomization, 100 nozzles at 1 L/min, and 85% plume delivery, the effective CCN flux at cloud base is approximately 2.1×10^{13} particles/s per vessel.

Estimated delivery. Nephele's rotary splash atomization system delivers approximately 2×10^{13} CCN/s to cloud base (range 1.7 to 2.3×10^{13} for plume delivery of 70 to 95%, daughter mass fraction 30-70%, and a CMD between 1.0 and 2.5 μm), figure 1. Whether this represents 10%, 50%, or even more than the pre-2020 per-ship cloud perturbation depends on how much of the pre-2020 signal came from new particle formation versus modification of existing particles, a question that published literature has not resolved and that bottom-up calculations cannot answer reliably. What is measurable is the cloud response itself. Phase 0 testing (section 14.1) will measure the actual albedo perturbation from a Nephele-equipped vessel and compare it directly to the pre-2020 ship track literature, removing the particle budget uncertainty. Regardless of performance, support from the IMO Net-Zero Fund will commercially incentivize improvements in this sector of engineering.

Bounding the pre-2020 comparison. Although a precise particle-for-particle comparison is not possible, we can bound the pre-2020 per-ship effective CCN rate at cloud base. The lower bound is set by primary particulate matter alone: 1.5 kg/hr of primary PM at CMD 100 nm gives approximately 5×10^{14} particles/s at the stack, of which perhaps 10 to 20% reach cloud base (consistent with the $\sim 10\%$ CCN/CN ratio measured by Hudson et al. 2000^[16]) with 50 to 70% activating as CCN, yielding approximately 3×10^{13} effective CCN/s. The upper bound adds a generous estimate of secondary sulfate nucleation: if 10% of the converted sulfate nucleates new particles (rather than condensing on existing ones) with 15 to 20% delivery and 60 to 70% activation, the total rises to approximately 10^{14} effective CCN/s. This roughly half-order-of-magnitude uncertainty range (3×10^{13} to 10^{14}) is shown as the shaded band in Figure 1, Panel B. Nephele's central estimate of 2.1×10^{13} falls at the lower edge of this band, meaning that with current hardware Nephele may deliver a comparable number of new CCN to what primary PM alone provided pre-2020, but likely falls short of the total pre-2020 cloud perturbation including secondary sulfate contributions. This bottom-up range is informed by measured ship emission factors^[15] and must be consistent with the observed $\sim 67\%$ cloud response reduction measured by Diamond and Boss (2025)[7]. This is the honest starting point. Phase 1 pilot testing (section 14.2) will measure whether Nephele's CCN flux produces a detectable cloud response.

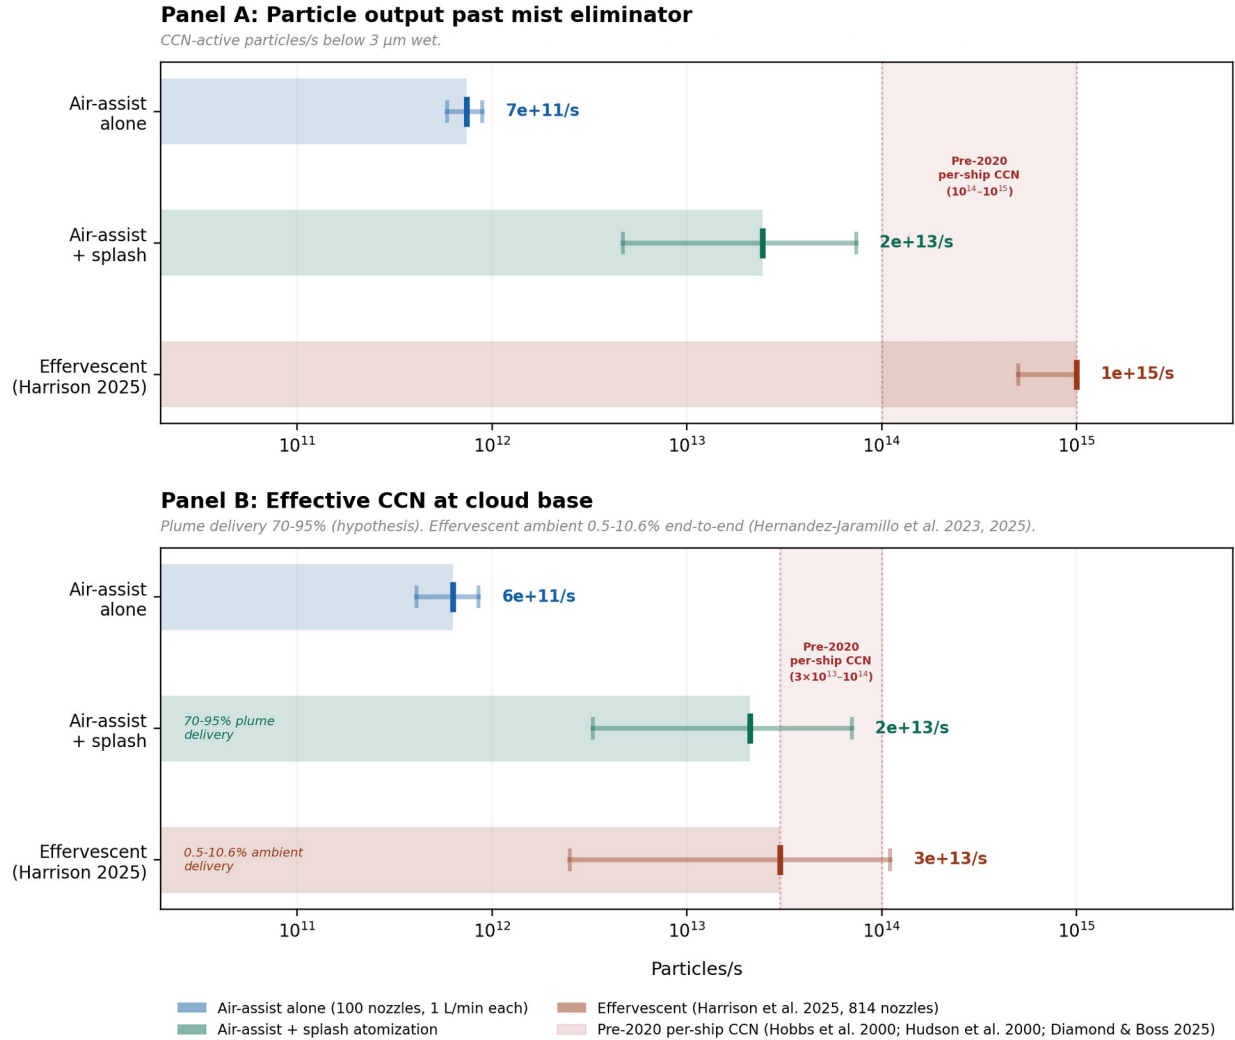


Figure 1. Panel A: CCN-active particle output past the 3 μm mist eliminator. Effervescent nozzles produce $\sim 40\times$ more useful particles because nearly all output is sub-100 nm. Rotary splash uncertainty reflects parameter ranges (mass fraction 30-70%, daughter count 8-25, daughter CMD 1.0-2.5 μm). Pre-2020 per-ship CCN band (10^{14} to 10^{15}) derived from Hobbs et al. (2000) total CN measurements with $\sim 10\%$ CCN/CN ratio (Hudson et al. 2000); upper end reflects dirtier fuel (2-3.5% S) than fleet average (2.7% S). Panel B: Effective CCN at cloud base after delivery losses. Nephele plume delivery (70-95%, hypothesis) preserves most output. Effervescent ambient delivery (0.5-10.6% end-to-end; Hernandez-Jaramillo et al. 2023, 2025) reflects passive MBL lofting and activation. Pre-2020 per-ship effective CCN band (3×10^{13} to 10^{14} , bottom-up estimate). Lower bound: primary PM only reaching cloud base. Upper bound: includes secondary sulfate nucleation. Informed by measured emission factors (Hobbs et al. 2000) and constrained by the observed $\sim 67\%$ cloud response reduction (Diamond and Boss 2025).

A note on plume dispersion. The plume spreads laterally as it rises, diluting injected particles across a widening cross-section. But pre-2020 ship tracks formed through exactly the same process. What matters is not absolute CCN concentration at cloud base but the perturbation relative to the clean background within the plume footprint. Ship tracks 10 to 50 km wide are routinely detected in satellite data^[27,44]. Dispersion is not a loss mechanism — it is the mechanism.

6.5 Current Hardware Limits and the Path to Parity

The 2×10^{13} CCN/s estimate depends on rotary splash atomization performing as modeled. Without splash atomization, air-assist nozzles alone deliver only a few percent of even the most conservative pre-2020 estimate. The rotary splash atomizer is unproven in continuous marine operation. The daughter droplet size distribution, splash mass fraction, and fragment count assumed here are based on published splash dynamics literature^[6,38,41], not on measurements from a Nephele prototype. Phase 0 hardware development (Section 14.1) will build and test the rotary splash atomizer alongside the measurement campaign.

Even with rotary splash atomization, off-the-shelf air-assist nozzles are unlikely to reach full parity with pre-2020 cloud perturbation levels. Closing the gap requires nozzle technology that puts substantially more output mass into the sub-3 μm wet range. Effervescent nozzles, de Laval flash atomization, further optimized rotor geometries, and other approaches are under active development. These exist in laboratory settings but have not been demonstrated for continuous unattended marine operation. Maturing them to fleet-ready reliability is an engineering program, not a physics problem.

This is precisely why IMO recognition of ship track restoration matters now, before the hardware reaches parity. If the MEPC recognizes verified radiative forcing from ship-integrated sea salt injection as an eligible activity under the Net-Zero Framework, it creates a direct commercial incentive for the maritime and industrial spray engineering sectors to develop nozzle hardware optimized for CCN production in marine exhaust environments. The market for such hardware does not currently exist because the regulatory category does not exist. Nephele's platform (mounting, verification instrumentation, fleet integration, regulatory positioning) is ready for improved nozzles the moment they are available. IMO recognition creates the market that funds the hardware development that closes the gap.

In the interim, 2×10^{13} clean CCN/s is not zero. Combined with the ship's own residual post-2020 combustion aerosol, a Nephele-equipped vessel produces a measurable, verifiable increase in cloud perturbation. Whether this is sufficient to produce a detectable ship track is the central question Phase 0 will answer. That measurement, not a bottom-up particle budget, is the test that matters.

7. Why Don't Scrubbers Already Do This?

They partially do. Every scrubber-equipped vessel is already emitting sub-micron aerosol that perturbs clouds at approximately one-third of pre-2020 levels.^[7] IMO 2020 capped sulfur in fuel, not at the scrubber outlet. There are simply fewer sulfate aerosols being generated by combustion. The scrubber's mist eliminator captures coarse washwater droplets but freely passes sub-micron particles; it functions as a particle size low-pass filter. Nephele exploits this same physics, deliberately injecting seawater spray below a mist eliminator to produce a controlled flux of CCN-active solution droplets.

8. Exhaust Chemistry and HCl

Sea salt aerosol in the marine atmosphere undergoes chloride depletion through reaction with sulfuric and nitric acid, releasing HCl. In the concentrated exhaust plume, residual SO_2 and NO_x could accelerate this. Three mitigations: (1) post-scrubber SO_2 is already substantially reduced to MARPOL compliance levels; (2) any HCl disperses into open atmosphere at levels orders of

magnitude below health concern; (3) chloride-depleted particles convert to sodium sulfate or sodium nitrate, both hygroscopic and still CCN-active, albeit at lower κ . The extent of in-plume chloride depletion is an open question for pilot testing.

9. Verification, Reporting, and Accountability

9.1 The Mist Eliminator Cassette as Verification Instrument

The replaceable cassette carries a time-integrated record of operation. During service, the mist eliminator captures coarse droplets rejected from the output stream. These deposits accumulate as a mixture of sea salt from the injected spray and sulfate absorbed from the exhaust gas. Laboratory analysis of a recovered cassette provides three independent verification measurements. First, total sodium chloride mass quantifies cumulative seawater throughput during the cassette's installed period, confirming injection volume independently of onboard flow meters. Second, the sulfate-to-sodium ratio in the deposits provides a time-integrated measurement of the exhaust stream's SO_2 concentration, since SO_2 dissolves into the liquid film on the mist eliminator wires and reacts to form sulfate in proportion to its gas-phase concentration. Third, total deposit mass combined with known spray composition constrains cumulative operating hours. Each cassette thus provides three verification layers: NaCl loading confirms injection throughput, sulfate assay verifies emissions compliance, and total mass constrains operating duration.

This has standalone value to the IMO MEPC for emissions verification of the 4,000+ scrubber fleet independent of ship track restoration.

9.2 Satellite and Shipboard Sensing Verification

The verification architecture is built in layers, calibrated during the pilot phase and simplified for fleet-scale deployment.

During Phase 1, tail aircraft flying through and above the plume will measure cloud droplet spectra, aerosol size distribution, and radiative fluxes downstream of injection vessels (Section 14.3). Simultaneously, shipboard instrumentation records injection parameters, meteorological conditions, and local aerosol concentrations. The aircraft data provides ground truth; the shipboard data provides the predictors. By statistically relating downstream cloud response to shipboard conditions across multiple transits and corridors, Phase 1 builds a transfer function: given these shipboard measurements, this is the expected cloud perturbation.

Once that transfer function is established, fleet-scale verification no longer requires tail aircraft on every vessel. The cassette confirms throughput (Section 9.1). Shipboard instrumentation confirms operating conditions. The Phase 1-derived statistical model converts these inputs to an estimated radiative forcing, which is then periodically validated by satellite observation and occasional dedicated aircraft campaigns.

The satellite tier operates at three levels. Primary: CERES broadband radiative flux measurements^[23,26] with coincident MODIS/VIIRS cloud retrievals.^[28] Secondary: machine-learning ship track detection cross-referenced with AIS vessel tracking.^[39] Tertiary: Diamond (2023) counterfactual statistical methods for corridor-scale forcing estimation.^[8]

The result is a verification stack that scales: cassette plus shipboard instrumentation on every vessel, satellite attribution continuously, and aircraft calibration periodically. Each layer is independently auditable.

9.3 Off Switch

Sub-micron sea salt aerosol in the MBL has atmospheric residence time of hours to days.^[29] Stop injection and cloud brightening ceases within 24 to 48 hours. No persistent atmospheric modification. No stratospheric loading. No multi-year commitment. Cessation authority is unconditional and retained by the MEPC, flag states, and port states.

10. What This Is Not

Not planetary-scale geoengineering. Ship track restoration returns a pre-existing forcing that we know existed and shipping itself produced for decades before being abruptly removed in 2020. We aim to restore that effect, not overreach it.

Not new science. The Twomey effect has been studied since 1977.^[36] Ship tracks since 1966. Salt aerosol CCN activation is well-characterised. Every mechanism in this proposal has been published and replicated. Engineering integration is minimized to an extreme, utilizing primarily off-the-shelf hardware.

Not permanent. Stop injecting and the effect vanishes within 48 hours. No other climate intervention has a faster off switch.

Not sulfur. We replace a harmful pollutant with a naturally occurring, ecologically benign substance that is a more effective CCN ($\kappa_{\text{NaCl}} \approx 1.28$ vs. $\kappa_{\text{sulfate}} \approx 0.61$), and far more predictable and controllable.^[31]

Not unregulated. Nephele proposes to operate exclusively under IMO MEPC oversight, on registered commercial vessels, in designated open-ocean corridors. Modules are programmatically deactivated in coastal waters, ports, and territorial seas.

Not encouraging dirty aerosol emissions. Nephele aims to develop the capacity to replace pre-2020 ship track strength entirely with clean sea salt, meaning we will not need combustion-generated aerosol. We welcome and encourage stricter fossil fuel emissions regulations.

11. Advantages Over Dedicated Cloud Brightening Proposals

Parameter	Dedicated MCB Systems	Nephele (Ship Track Restoration)
Vessel fleet	Purpose-built spray ships	4,000+ existing scrubber vessels
Spray system	Precision sub- μm nozzles	Air-assist nozzles + rotary splash atomizer + mist eliminator
Delivery mechanism	Ambient MBL lofting (1 to 10% efficiency)	Exhaust plume (70 to 95% efficiency, hypothesized)
Particle state at cloud base	Dry crystal (must deliquesce + activate)	Effectively pre-activated solution droplet
Power source	Dedicated generators	Existing ship engine
Seawater supply	Dedicated pumps/intakes	Existing scrubber circuit
Deployment timeline	Years to build fleet	Months to retrofit

Per-unit cost	\$10M+ per vessel	~\$250K per module
Regulatory pathway	Undefined	IMO MEPC oversight pathway defined
Verification	Satellite + in-situ sensing	Cassette + satellite + in-situ sensing

12. Regulatory Pathway

12.1 IMO Net-Zero Framework

The IMO Net-Zero Framework was approved in principle at MEPC 83 in April 2025, with formal adoption adjourned at the extraordinary session in October 2025 (vote: 57 to 49 to defer). Entry into force is expected no earlier than March 2028.[18] Revenue flows into the IMO Net-Zero Fund, mandated to reward zero and near-zero emission technologies.[19]

12.2 ZNZ Definition

Regulation 31.19 of draft revised MARPOL Annex VI defines ZNZs as "zero or near-zero GHG emission technologies, fuels and/or energy sources." [20] The definition explicitly includes "technologies" as a standalone category. The ZNZ guidelines remain under development as of April 2026; ten new supporting guidelines were identified at MEPC 83.[18] The guidelines gap is Nephele's window.

12.3 Legal basis for ZNZ eligibility

The standard ZNZ threshold under Regulation 39.1 is defined as a GHG fuel intensity not exceeding 19.0 gCO₂eq/MJ on a well-to-wake basis, tightening to 14.0 gCO₂eq/MJ from 1 January 2035. Sea salt aerosol injection does not consume fuel and produces no GHG emissions. It has no GFI value. A strict reading would exclude it, not because it fails the test, but because the metric does not apply.

The critical provision is the final sentence of Regulation 39.1: "Notwithstanding, the Committee may approve additional ZNZs taking into account guidelines developed by the Organization." (18)

This sentence is the legal basis for Nephele's inclusion. It grants the MEPC explicit discretionary authority to recognise ZNZ technologies that fall outside the standard fuel-intensity framework. The accompanying footnote 114 in the Annex directs this authority to guidelines to be developed — guidelines whose scope and criteria have not yet been written. Those guidelines are being developed ahead of MEPC 85 in late 2026, and the drafting window is open now.

Nephele's regulatory objective is precise: ensure that the ZNZ guidelines developed under footnote 114 include a pathway for ship-installed technologies producing verified net negative radiative forcing. This does not require amending draft MARPOL Annex VI. It requires only that the guidelines — which the Committee has already tasked itself to produce — are written to encompass this category.

The substantive case for inclusion rests on four grounds. First, the Net-Zero Framework measures climate impact, and Nephele produces a verified, quantifiable, negative climate impact — the

framework's stated purpose. Second, the IMO is itself the author of the cooling gap Nephele fills: IMO 2020 eliminated decades of inadvertent ship track forcing; the Net-Zero Fund is the appropriate vehicle to restore it. Third, precedent for non-fuel ZNZ recognition already exists within the framework, which explicitly lists "technologies" and "energy sources" alongside fuels as eligible categories. Fourth, the restoration framing is legally and politically distinct from conventional geoengineering: Nephele returns a pre-existing, well-documented forcing that shipping produced for over a century and abruptly removed — it does not propose a novel planetary intervention.

12.4 What We Need from the MEPC

Three options exist, in order of preference, all achievable within the current guidelines process:

Option A (preferred): The ZNZ guidelines developed under Regulation 39.1 include an explicit pathway for ship-installed technologies producing verified net negative radiative forcing, with the W/m^2 to CO_2eq equivalence methodology defined as part of the same guidelines package. This is the cleanest outcome and the one Nephele's regulatory engagement is directed toward.

Option B: The guidelines define ZNZ eligibility broadly enough — consistent with the "technologies" language already in Regulation 39.1 — that any ship-installed technology producing verified net climate benefit qualifies, without requiring a named category. Nephele would then seek formal recognition through the standard guidelines application process.

Option C (fallback): A separate disbursement category is created under the Fund's climate protection and environmental resilience mandate (Regulation 40), distinct from the standard ZNZ reward mechanism but drawing from the same Fund. This would require more procedural development but remains within the existing legal framework.

None of these options requires a full reworking of draft MARPOL Annex VI; the text is settled, the guidelines are not.

12.5 MEPC Oversight as Strategic Position

The proposed hardware is not proprietary. Satellite data is public. Our moat is regulatory: the institutional relationship with the MEPC, the approved verification methodology, the track record of auditable operations, and the trusted-operator status from helping write the rules and operating transparently under them.

13. Revenue Model

13.1 Primary revenue: IMO Net-Zero Fund Disbursements

The IMO Net-Zero Framework, once formally adopted and operational from 2028, will collect revenue through mandatory GHG pricing contributions from non-compliant vessels and disburse those funds as rewards to ships operating zero or near-zero emission technologies. The fund is projected to generate \$10–15 billion annually in its initial years of operation[35].

The reward mechanism operates on a per-tonne-of-CO₂eq basis, with the specific rate to be defined by the Committee no later than 1 March 2027 under Regulation 39.3. One additional parameter must be established before Nephele can operate within this mechanism: a formally adopted methodology for converting verified radiative forcing, measured in W/m², into CO₂eq avoided. This conversion is scientifically grounded in the same peer-reviewed literature that quantifies the warming effect of IMO 2020, and developing this equivalence methodology is a core objective of Nephele's MEPC engagement.

Once both parameters are set — the reward rate and the W/m² conversion — the payment flow is straightforward. Verified radiative forcing from each participating vessel, measured via the satellite attribution pipeline described in Section 9, is converted to CO₂eq avoided, multiplied by the agreed reward rate, and disbursed by the Fund on an annual reporting cycle. Nephele invoices the Fund for verified forcing delivered across its fleet, retains a contracted margin, and remits the balance to partnered shipping operators as compensation for carrying the injection equipment — offsetting a portion of their own GHG compliance costs under the same framework.

What is already knowable is the structure: the cost base is largely fixed once the verification infrastructure is built, per-vessel economics improve as installation costs are recovered against recurring fund disbursements over time, and overhead is spread across a growing fleet. Per-vessel projections will be developed as the reward rate and conversion methodology are defined through the guidelines process.

13.2 Interim funding

NZF disbursements begin no earlier than 2029. The path to that point is manageable: Nephele is not capital-intensive in its early years. The primary costs are people and the Phase 0/1 measurement campaigns. This means Nephele can be selective about external capital, prioritising grant funding and strategic partnerships over venture investment during the period when regulatory positioning matters most.

On the grant side, ARIA's Exploring Climate Cooling programme is the strongest near-term fit: ARIA has committed £56.8M across 22 climate cooling programmes, the scientific framing of ship track restoration maps cleanly onto their mandate, and Nephele's open-ocean regulatory approach addresses the community engagement concerns that have complicated onshore MCB research

elsewhere. Innovate UK and philanthropic climate foundations represent additional near-term sources. Aggregate realistic potential from non-dilutive grant funding is \$1M+.

The second pillar is pre-commercial shipping partnerships. Pilot agreements with scrubber-equipped operators before 2029 generate no fund revenue directly, but they produce something more valuable at this stage: the installed base, measurement dataset, and operational track record that makes Nephele's commercial position defensible from day one of fund disbursements. A partner who has run a Nephele module for two years before the fund opens is not a prospect — they are a paying customer waiting for the mechanism to activate.

The verification infrastructure and satellite attribution pipeline, once built for the pilot fleet, scale across additional vessels at low marginal cost. Hardware per vessel is approximately \$250,000 (Section 5.5). The scrubber industry established precedent ahead of IMO 2020: capex-sharing clauses, profit-share arrangements, and third-party leasing structures emerged as the payback calculus became clear. The Nephele module is removable, cassette-based, and independent of the vessel's existing financing. The \$250,000 upfront is financeable for any operator already facing GHG compliance costs under the same framework.

13.3 Compliance verification: an independent revenue stream

The Nephele cassette, described in Section 9.1, carries a tamper-resistant physical record of scrubber operation: NaCl mass quantifies cumulative injection throughput, and a sulfate-to-sodium ratio independently verifies the exhaust stream's SO₂ concentration during the cassette's operating period. This is relevant not only to Nephele's ship track restoration activity but to the broader challenge of scrubber emissions verification across the global fleet.

MARPOL Annex VI requires scrubber-equipped vessels to demonstrate compliance with sulfur emission limits, but current verification relies primarily on onboard electronic logging with periodic port state control inspections, a system with known monitoring gaps. The cassette provides a physical, chemistry-based verification record that is independent of onboard electronics, tamper-evident by nature, and recoverable post-voyage. This has standalone value to classification societies, flag states, and the IMO as a verification instrument for the 4,000+ scrubber-equipped vessels operating today, entirely independent of ship track restoration.

The commercial model for this revenue stream can be developed separately, but the outline is straightforward: Nephele manufactures and supplies easily-removable, tamper-resistant cassettes to scrubber operators on a subscription basis, retrieves and analyzes them on a defined cycle, and provides verified emissions records to operators, classification societies, or flag state authorities as required. This business generates revenue from the moment cassettes are in operation, well before fund disbursements begin, and is fundable through the same grant and partnership pathways as the core ship track restoration activity.

14. Pilot Program

14.1 Phase 0: Measurement Campaign and Initial Development

Phase 0 has two parallel tracks. The first is a measurement campaign to quantify CCN concentration in post-2020 scrubber-equipped ship plumes and compare to indirect measurements and statistical estimates in literature. Plume measurements should also aim to quantify plume dispersion and relative humidity changes with altitude. This data does not exist publicly in peer-reviewed literature.

The second track is initial hardware development: bench-scale nozzle and rotary atomizer testing, mist eliminator characterization with seawater spray, cassette deposit characterization (NaCl throughput quantification and sulfate uptake rates under representative exhaust SO₂ concentrations), and prototype module design. Because the nozzle strategy is deliberately flexible (Section 5.3), hardware development proceeds in parallel with the measurement campaign. The measurement campaign tells us how large the gap is. Early engineering development ensures we have a module ready to fill as much of that gap as possible.

14.2 Phase 1: Pilot Deployment

Modules installed on 1 to 10 scrubber-equipped cargo vessels operating regular routes through marine stratocumulus regions.

Phase 1 serves three purposes. First, it measures whether Nephele's CCN flux produces a detectable cloud response; this is the central empirical test. Second, it calibrates the verification architecture described in Section 9.2: tail aircraft measure downstream cloud response while shipboard instrumentation records operating conditions, building the statistical transfer function that allows fleet-scale verification without aircraft tailing every vessel. Third, it establishes a track record in real-world operations: engineering, installation, testing, and regulation applied at a meaningful scale.

Measurement tiers. Tier 1 runs on every vessel, continuously: optical particle counter above the mist eliminator, differential pressure sensing, shipboard weather station, GPS, injection parameters (flow rate, pressure), ceilometer, AIS. Tier 2 overlays periodically with dedicated research aircraft flying through and above the plume, measuring aerosol size distribution, relative humidity, cloud droplet spectra, drizzle rates, and radiative fluxes. Campaigns of 2 to 4 weeks, repeated 2 to 3 times across different corridors and seasons. The aircraft data provides ground truth for the cloud response; the shipboard data provides the predictors. By correlating the two across multiple transits, Phase 1 establishes which shipboard conditions produce which cloud outcomes — the basis for the adaptive injection algorithm and the fleet-scale verification model (Section 9.2).

Stopping criteria. Stopping criteria will be defined before the pilot begins, preferably in close collaboration with the IMO MEPC. We suggest some guidelines here that may be reworked: Injection is halted if: (1) net cloud darkening is observed over >10 independent transits; (2) drizzle rates exceed the statistically defined baseline by >2× over 5+ transits; (3) the mist eliminator passes >5% by number of particles >3 μm wet; (4) the module interferes with scrubber MARPOL compliance; (5) the MEPC, flag state, or port state directs cessation. All criteria and triggering thresholds will be published in advance. Tail aircraft may observe adverse effects in real time and halt injection mid-transit. All results should be publicly published.

Adaptive injection control. Each vessel's Tier 1 instrumentation feeds an injection decision algorithm that activates in designated corridors when shipboard conditions match the brightening-favorable profiles identified during the pilot, and deactivates when conditions are unfavorable or when the vessel exits the approved corridor. The algorithm is conservative by design: injection

defaults to off and activates only when conditions are positively identified as favorable. Phase 1 testing with Tier 2 tail aircraft may allow real-time tuning of this algorithm.

15. Proposed Regulatory Framework

Nephele will not deploy without regulatory backing.

The regulatory framework should cover: equipment type-approval by classification societies; designation of approved operating corridors; monitoring, reporting, and verification requirements; maximum injection rates in terms of CCN concentration at the ship exhaust, and minimum vessel separation; absolute cessation authority retained by MEPC, flag states, and port states (unconditional, not subject to appeal); liability and insurance requirements; and a mandatory periodic review cycle.

16. Risks and Mitigations

Risk	Severity	Mitigation
Net cloud darkening from drizzle induction	High	Stopping criteria; aircraft drizzle measurements; mist eliminator rejects >3 μm ; open-ocean only
Rotary splash atomization underperforms	Medium	Phase 0 bench testing; rotary atomizer is modular and replaceable; nozzle-agnostic platform allows technology swap
Insufficient CCN production	Medium	Phase 0 quantifies gap; plume delivery multiplier; nozzle upgrades as technology matures
IMO/MEPC rejects ZNZ classification	High	Early delegation engagement; multiple regulatory options (A/B/C); NZF mandate is broad
Political "geoengineering" backlash	High	MEPC oversight; open data; 48-hour reversibility (to be demonstrated in Phase 0/1); "restoration not engineering" framing; open-ocean only;
Plume delivery assumptions wrong	Medium	Phase 0/1 measurements directly test delivery efficiency; stopping criteria if cloud response insufficient
Mist eliminator / rotor blade fouling	Low	Replaceable cassette; dP monitoring; routine port-call swap; corrosion-resistant materials
Moral hazard (delay decarbonization)	Medium	NaCl not combustion-linked; system capacity can exceed pre-2020 levels with future deployment of improved nozzles; we welcome stricter fuel standards

17. Conclusion

The IMO 2020 sulfur cap removed a harmful pollutant from shipping emissions. It also removed the aerosol that had been seeding brighter marine clouds for over a century. The clouds dimmed. The ocean warmed.

Nephele puts the reflectivity back. Pure seawater aerosol, delivered through the ship's own exhaust plume. The plume keeps droplets in solution from exhaust stack to cloud base: no droplet drying, no deliquescence, minimal lofting losses. Ship plumes are the most efficient CCN delivery method on the planet, and 4,000+ of them are already at sea and ready for retrofit with Nephele systems. The ship supplies the seawater, the energy, and the ride.

Current hardware delivers approximately 2×10^{13} clean CCN/s to cloud base. Whether this produces a detectable ship track is an empirical question Phase 0 will answer. Full parity with pre-2020 levels will require nozzle technology optimized for marine CCN production — hardware that does not yet exist as a commercial product, because the market for it does not yet exist.

IMO recognition creates that market. The Net-Zero Framework provides the vehicle. The physics are established. The verification infrastructure exists. Nephele's platform is ready for better nozzles the moment they arrive.

The first step is a measurement campaign: characterize post-2020 scrubber plumes and quantify the CCN deficit. This data does not exist in the peer-reviewed literature. Nephele will produce it.

Ship track restoration is the deliberate, clean, verifiable return of forcing that shipping itself produced for decades and abruptly removed. The regulatory framework that recognizes it will also fund the hardware that perfects it.

References

- [1] ACS Industries. The Engineered Mist Eliminator. Technical Bulletin.
- [2] Brendecke, J., et al. (2022). Maritime Aerosol and CCN Profiles Over Eastern North Pacific During MAGIC. *Earth and Space Science*, 9.
- [3] Bretherton, C. S., et al. (2010). VOCALS-REx: Observations of SE Pacific stratocumulus. *Atmos. Chem. Phys.*, 10, 10639. [MBL RH never observed below 45% in SE Pacific Sc region; typical sub-cloud RH 80-95%]
- [4] Cadeddu, M. P., et al. (2023). Boundary layer moisture variability at the ARM Eastern North Atlantic observatory. *Atmos. Chem. Phys.*, 23, 3453.
- [5] Carbon Brief (2023). Analysis: How low-sulphur shipping rules are affecting global warming.
- [6] Cossali, G. E., Coghe, A. & Marengo, M. (1997). The impact of a single drop on a wetted solid surface. *Exp. Fluids*, 22, 463. [Splash mechanisms on wetted surfaces; thin films promote splashing]
- [7] Diamond, M. S. & Boss, L. F. (2025). Conflict constrains cloud sensitivity to stricter marine pollution regulations. *Atmos. Chem. Phys.*, 25, 16401.
- [8] Diamond, M. S. (2023). Detection of large-scale cloud microphysical changes within a major shipping corridor. *Atmos. Chem. Phys.*, 23, 8259.
- [9] Koch-Glitsch. Demister Mist Eliminator. <https://www.koch-glitsch.com/product-lines/mist-elimination/demister>
- [10] Gettelman, A., et al. (2024). Has Reducing Ship Emissions Brought Forward Global Warming? *Geophys. Res. Lett.*, 51, e2024GL109077.
- [11] Guildenbecher, D. R., et al. (2009). Secondary atomization. *Exp. Fluids*, 46.
- [12] Hansen, J. E., Kharecha, P., Morgan, D. & Vest, J. (2026). Super El Nino? Super Warming is the Main Issue. 20 March 2026. [Climate sensitivity 4-5°C for doubled CO₂; reduced aerosol cooling including shipping driving accelerated SST warming]
- [13] Harrison, L. P., et al. (2025). Effervescent nozzle design for MCB experimentation. *Environ. Sci.: Atmos. (RSC)*.
- [14] Hernandez-Jaramillo, D. C., et al. (2023/2025). A review of plume dispersion and measurement techniques applicable to MCB. *Frontiers in Marine Science*.
- [15] Hobbs, P. V., et al. (2000). Emissions from ships with respect to their effects on clouds. *J. Atmos. Sci.*, 57, 2570-2590.
- [16] Hudson, J. G., Garrett, T. J., Hobbs, P. V., Strader, S. R., Xie, Y. & Yum, S. S. (2000). Cloud condensation nuclei and ship tracks. *J. Atmos. Sci.*, 57, 2696-2707. [~10% CCN/CN ratio for HFO-burning ships at 0.2% supersaturation]
- [17] IFRF Handbook. What are main characteristics of a black liquor recovery boiler? [Splash plate nozzle as most common type; mean droplet size 2-3 mm]
- [18] IMO (2025). MEPC 83 Press Briefing. Korean Register, IMO News Flash MEPC 83.
- [19] IMO (2025). The IMO Net-Zero Framework FAQs.
- [20] IMO (2025). Resolution MEPC.403(83): Draft Amendments to MARPOL Annex VI.

- [21] Jenkins, A. K. L. & Forster, P. M. (2013). The inclusion of water with the injected aerosol reduces the simulated effectiveness of MCB. *Atmos. Sci. Lett.*, 14, 164. [Evaporative cooling suppresses plume lofting]
- [22] Jordan, G., et al. (2024). IMO2020 Regulations Accelerate Global Warming by up to 3 Years. *Earth's Future*, 12, e2024EF005011.
- [23] Kato, S., et al. (2022). CERES Cloud Radiative Swath Data Product. *J. Atmos. Oceanic Technol.*, 39.
- [24] Latham, J., et al. (2012). Marine cloud brightening. *Phil. Trans. R. Soc. A*, 370, 4217.
- [25] Lefebvre, A. H. & McDonnell, V. G. (2017). *Atomization and Sprays*, 2nd ed. CRC Press.
- [26] Loeb, N. G., et al. (2018). CERES EBAF TOA Edition-4.0. *J. Climate*, 31, 895.
- [27] Manshausen, P., et al. (2022). Invisible ship tracks show large cloud sensitivity to aerosol. *Nature*, 610, 101.
- [28] Minnis, P., et al. (2021). CERES MODIS cloud product retrievals. *IEEE Trans. Geosci. Remote Sens.*, 59.
- [29] Murphy, D. M., et al. (2019). Distribution of sea-salt aerosol in the global troposphere. *Atmos. Chem. Phys.*, 19.
- [30] Neukermans, A., et al. (2014). Sub-micrometer salt aerosol production for MCB. *Atmos. Res.*, 142.
- [31] Petters, M. D. & Kreidenweis, S. M. (2007). Single parameter representation of hygroscopic growth and CCN activity. *Atmos. Chem. Phys.*, 7.
- [32] Raju, V., et al. (2021). Kraft recovery boiler operation with splash plate and/or beer can nozzles. *TAPPI Journal*. [CFD modeling of splash plate nozzle performance]
- [33] Ramana, M. V., et al. (2016). CCN concentrations influenced by ship emitted aerosol plumes. *Scientific Reports*, 6, 30416.
- [34] ScienceDirect (2008). Mist Eliminator overview. *Gas-Liquid And Liquid-Liquid Separators*.
- [35] TESS Forum (2025). Equity in Motion: Revenue Disbursement Under the IMO Net-Zero Framework.
- [36] Twomey, S. (1977). The influence of pollution on the shortwave albedo of clouds. *J. Atmos. Sci.*, 34.
- [37] Vander Wal, R. L., Berger, G. M. & Mozes, S. D. (2006). Droplets splashing upon films of the same fluid of various depths. *Exp. Fluids*, 40, 53. [Thin films promote splashing; thick films inhibit it]
- [38] Vimalraj, S. R. & Vaikuntanathan, V. (2024). Splash droplet size in drop impact on solid surfaces. [Daughter droplet diameter scales as $We^{-1/4}$]
- [39] Watson-Parris, D., et al. (2022). Shipping regulations lead to large reduction in cloud perturbations. *Proc. Natl. Acad. Sci.*, 119.
- [40] Wood, R., et al. (2024). Physical science research needed to evaluate MCB. *Science Advances*.
- [41] Yarin, A. L. (2006). Drop impact dynamics: splashing, spreading, receding, bouncing. *Annu. Rev. Fluid Mech.*, 38, 159.
- [42] Yoshioka, M., et al. (2024). Warming effects of reduced sulfur emissions from shipping. *Atmos. Chem. Phys.*, 24, 13681.
- [43] Yuan, T., et al. (2024). Abrupt reduction in shipping emission as an inadvertent geoengineering termination shock. *Commun. Earth Environ.*, 5, 281.

[44] Yuan, T., et al. (2025). Detectable ship tracks account for just 5% of aerosol indirect forcing. *Commun. Earth Environ.*, 6, 899.

Contact • nephele.earth • info@nephele.earth
© 2026 Nephele. All rights reserved.