

AETHER v0.46 technical supplement

Noah Hicks | Technical supplement v0.46 | September 2026

Companion to: *AETHER: Atmospheric Engineering Through High-Energy Removal*

Status: Internal technical inventory; not external peer review, validation, or engineering certification

Purpose

This supplement preserves the useful work of the legacy model suite without presenting it as a single scientific result. It is an inventory of screening tools, their intended role, and the boundary on their use after the September 2026 correction. The principal paper contains the argument; this document is where an interested reviewer can find the model architecture, calculations worth reproducing, and the limitations that prevent promotion to a forecast.

Generated tables, source notes, scripts, and older rendered manuscripts remain in the repository for reproducibility and historical comparison. Their presence does not mean every output is current evidence.

Accounting definitions

The suite previously mixed quantities that need separate treatment. In v0.46:

- **Gross capture/removal** is the CO₂ sent into a removal process.
- **Retained gross** is captured CO₂ multiplied by a pathway-specific retention term.
- **Lifecycle debit** is reported separately from retained gross in simplified annual ledgers.
- **Net-after-lifecycle result** is the signed physical-accounting result after the stated lifecycle debit; it is not an atmospheric trajectory, and a negative value is a reported burden.
- **Creditable removal** applies a stated measurement, reversal, leakage, and invalidation treatment after physical accounting; only this layer is floored at zero.
- **Net climate result** requires time-indexed emissions, removal, carbon-cycle response, and non-CO₂ treatment. It cannot be inferred from a scalar credit ledger.

Electricity, temperature-specific useful heat, and chemical energy must remain separate inputs. A scalar lifecycle CO₂e debit is not a substitute for species-resolved and time-resolved climate accounting.

Model inventory and use status

Model family	What it can do	v0.46 status and required interpretation
Mass-energy screens	Convert stated GJ/tCO ₂ , flows, and storage choices into annual energy and power orders of magnitude.	Retained as transparent arithmetic. Inputs are scenario assumptions unless directly sourced. Do not treat thermodynamic lower bounds as deployable process performance.
Hardware, air-throughput, and materials screens	Make contactor, pressure-drop, sorbent, steel, cement, copper, and plant-count implications visible.	Constraint inventory only. It is not a factory-rate, supply-chain, or regional deployment study.
Pathway portfolio and storage screens	Compare pathway assumptions and represent storage-route, injection, and durability terms.	Not additive without shared-resource constraints. Storage capacity is not injection rate, permitting capacity, monitoring capacity, or durable liability coverage.
Regional dispatch and colocation screen	Test named synthetic-day supply/load shapes and storage bookkeeping.	Corrected cyclic-storage diagnostic only. It is not 8760-hour, weather-resolved, transmission-resolved, or a regional feasibility result.
Carbon-cycle and climate screens	Test transparent pulse-response and forcing relationships.	Absolute projections are quarantined. Both the legacy fixed baseline and the candidate published-reference/Joos anomaly hybrid fail the required baseline boundary; neither supports target dates or validated temperature claims.
Lifecycle and MRV screens	Show the consequences of explicit retention, lifecycle, and buffer terms.	Corrected ledger separates retained gross and lifecycle debit. Terms remain provisional until pathway-specific LCA and measurement evidence replace them.
Cost and learning screens	Identify which cost components and learning assumptions control a scenario.	A partially coupled screen uses one-year-lagged cumulative realized removal, not planned capacity. Its 0.00204 Gt initial learning state is an explicit unsourced screening proxy, not a factual history; cost floors are assumptions.

Model family	What it can do	v0.46 status and required interpretation
Robotics and AI screens	Translate automation claims into task hours, supervision, maintenance, production, and availability questions.	Scenario and research-design tools. Company claims, AI forecasts, and assumed productivity are not field evidence.
Monte Carlo and correlated families	Expose break-even surfaces and failure modes under documented sampled assumptions.	Hand-set scenario distributions, not probabilities, forecasts, confidence intervals, or expected outcomes. Correlated families also change marginal distributions.
Integrated feasibility screen	Jointly constrain its own planned capacity, energy, robots, storage, budget, emissions, and rebound terms.	Partially coupled screening only. It does not yet use a common compatible case with upgraded regional dispatch, lifecycle/MRV, or field-productivity screens.

Corrected regional storage diagnostic

The earlier representative-day dispatch calculation initialized a battery at 50% state of charge and annualized the day without an equal terminal state. That effectively supplied initial stored energy repeatedly. The correction solves for a cyclic one-day storage boundary using the same rounded 24-hour synthetic profiles and annualizes the resulting day 365 times.

Legacy case	Earlier supported removal (GtCO2/year)	Cyclic-storage diagnostic (GtCO2/year)
Market reference	15.46	9.87
Dedicated corridors	47.93	33.47
Firm co-located backbone	64.26	42.72
Upper-tail AI/energy abundance	121.97	98.40
Fragmented nonadditional grid	10.19	3.62

The correction is a conservation requirement, not an argument that the five figures are realistic annual capacities. The profile, hardware, demand, location, and buildout assumptions remain screens. A real regional study needs weather years, load, transmission, curtailment, outage, maintenance, storage-duration, and additionality treatment.

Regional 1 MtCO2/year mechanism benchmark

The new regional reference is a matched analytical test anchored to NETL's solvent DAC Case 1 (NETL, 2022). It fixes auxiliary electricity at 0.533129 MWh/tCO2 and the calciner fuel-input-equivalent proxy at 6.846332 GJ_HHV/tCO2. The latter derives from natural-gas HHV input, not useful heat delivered at the required temperature. Replacing that input with low-carbon thermal supply requires separate conversion and integration evidence; no efficiency is invented here. The case is not a construction recommendation, field result, or general DAC performance estimate.

Result (tCO2/year unless stated)	Ordinary	Automation-assisted
Gross captured	850,000.0	900,000.0
Gross stored	847,451.7	897,301.8
Project emissions (tCO2e)	45,090.7	47,743.1
Retained physical-accounting result	843,214.4	892,815.3
Signed net after project emissions (tCO2e)	798,123.7	845,072.2
Risk-adjusted credits	758,696.4	803,325.6

Both cases share the stated removal process, energy and storage boundary, and wider cost/funding assumptions. The automation-assisted branch changes only uptime (0.85 to 0.90), task hours (660,000 to 345,000/year), and its \$12 million automation-system cost at the same \$85/hour task cost. It produces 50,000 additional gross tonnes/year because of the uptime assumption. Risk-adjusted credits use an EPA-informed storage/MRV boundary (U.S. EPA, 2026); they are not an independent climate forecast.

The annual emissions proxy scales with throughput. It does not independently resolve fixed embodied emissions, construction timing, decommissioning, or continuing monitoring and remediation during shutdown. These omissions matter for low-output and failed projects. Its allowance/allocation energy ledger is not generation or curtailment data. The case requires a full process-specific lifecycle and post-closure analysis before deployment claims.

Carbon and climate method boundary

The original concentration calculation attached future pulses to a fixed contemporary concentration. A zero-future-emissions test that never changes concentration exposes the missing reservoir initialization. Its absolute trajectories are withdrawn from evidentiary use.

The intended replacement draws annual CO2 concentration and total-CO2-emissions trajectories from the RCMIP v5.1.0 SSP2-4.5 reference dataset (Nicholls & Lewis, 2021; Nicholls et al.,

2020). Its first implementation applied a Joos impulse-response perturbation to the future difference between a named AETHER scenario and that reference, then added the published reference concentration (Joos et al., 2013). Matched no-AETHER controls are required when emissions assumptions differ.

The source provenance is verified, but the first absolute-output method is quarantined: its zero-future-emissions diagnostic fell and then rose from 426.582160 to 433.368307 ppm under a large off-reference anomaly. It is therefore not a trustworthy zero-emissions-commitment baseline. Future replacement work must resolve reservoir initialization before it is called a hybrid conditional baseline. Aggregate historical non-CO2 forcing remains synthetic in the legacy emulator. Temperature claims await paired, species-emissions climate work.

Cost and finance boundary

The revised learning screen uses a one-year lagged, cumulative realized-removal state. In the reference screen its 2026--27 state is 0.00204 Gt and its displayed \$500/t cost is a named screening input; 2028 begins from 0.15963 Gt of modeled realized removal. The initial state is not a claim about actual cumulative global deployment. It prevents an unbuilt planned ramp from manufacturing learning progress in the model.

For the 100 Gt/year low-carbon-mixed-power lifecycle screen, 95.864795 GtCO₂/year is retained gross; the separately reported lifecycle debit is 12.296750 GtCO₂e/year; the scalar net-accounting proxy is 83.568044 GtCO₂e/year; and provisional MRV credit is 65.086367 GtCO₂e/year. The old value of 83.915860 is superseded. None of these is a time- or species-resolved climate flow.

Costs are not harmonized, audited future prices or tariff recommendations. The service architecture needs a stock-flow model that separates operating outlays, capital, legacy drawdown, monitoring, reversal reserves, and a shrinking emissions-fee base. Any financial result should state who pays, when, what happens under under-delivery, and whether public balance sheets absorb the tail risk.

Worked accounting example

Take 1,000 gross tonnes entering the regional capture process. The benchmark assumes 99.8% transport delivery, 99.9% injection acceptance, and 99.5% retention over its stated accounting treatment:

$$\text{Stored} = 1,000 \times 0.998 \times 0.999 = 997.002 \text{ tCO}_2$$

$$\text{Retained} = 997.002 \times 0.995 = 992.01699 \text{ tCO}_2$$

The separate project-emissions proxy combines electricity, the thermal-energy scenario and the other-lifecycle placeholder. Per gross tonne it is:

$$\text{Debit} = (0.5331294234 \times 0.020) + (6.8463317661 \times 0.004) + 0.015$$

The products have units of tCO₂e per gross tCO₂. For 1,000 tonnes, the debit is about 53.04792 tCO₂e. Thus the signed net-accounting result is 938.96907 tCO₂e. Only credit eligibility is floored at zero; the 2% measurement discount and 3% risk buffer then give:

$$\text{Credits} = \max(938.96907, 0) \times 0.98 \times 0.97 = 892.58400 \text{ tCO}_2\text{e}$$

Multiplying this worked quantity by 850 gives the ordinary case's 758,696.4 annual risk-adjusted credits. If emissions instead exceeded retention, net benefit would be negative and issued credits zero. Retention cannot shrink the emissions debit. None of the assumed percentages is established by a site study or an EPA endorsement.

The cash identity is independent: sources = current-load base x settlement rate + legacy funding; balance = sources - annual uses. The regional model caps throughput at the smallest physical or budget limit and stops operation/credit issuance when verification checks fail. A budget below fixed obligations can still produce a deficit at zero throughput: stopping capture does not erase liabilities.

Reproduction and reviewer tasks

The immediate reviewer tasks are bounded:

- Reproduce the cyclic storage calculation and test conservation across all dispatch cases.
- Inspect the paired RCMIP/Joos baseline and verify that non-removal assumptions match between intervention and control.
- Reproduce the 1 MtCO₂/year ordinary-versus-automation regional case from source-linked inputs.
- Challenge the lifecycle/MRV ledger, especially the separation of physical removal, emissions debit, and credit buffer.
- Audit the finance model for solvency under legacy drawdown, post-closure liability, and falling fee revenue.