0G Validator Layer Scope 2 Emissions Methodology – Mainnet (MiCA-Aligned)

1) Scope & Purpose

This document defines a MiCA-aligned methodology to estimate Scope 2 emissions (electricity) for the validator layer of 0G on mainnet, with a testnet case study based on real validator locations and power profiles captured earlier in this thread. It mirrors the format of the storage methodology and is designed for repeatable reporting and regulator review.

Boundary: All validator nodes running the software, including those in the active set (participating in consensus) and those not in the active set (e.g., staked candidates or inactive validators maintaining sync). Storage, DA, compute, and serving layers are out of scope here.

2) Validator Hardware

Component	Mainnet	Testnet
Memory	64 GB	64 GB
CPU	8 cores	8 cores
Disk	1 TB NVMe SSD	4 TB NVMe SSD
Bandwidth	100 Mbps up/down	100 Mbps up/down

Power model:

- Testnet: CPU+RAM+NIC+misc ≈ 188 W IT; NVMe (4 TB) ≈ 11 W IT → 199 W IT. PUE ≈ 1.4 → ~280 W wall.
- Mainnet: 1 TB NVMe ≈ 3 W IT vs 11 W for 4 TB → −8 W IT. With PUE ≈ 1.4 → −11.2 W wall. Result: ~269 W wall.

Note: Empirical wall power measurement will be conducted using our own validator hardware post-mainnet launch to validate these modeled values.

3) Calculation Method

Let N(d) = online validators (active + powered standby + powered inactive/candidate) on day d.

1. Total network power (W):

P_net(d) = N(d) * P_node

2. Daily energy (kWh):

E_day(d) = (P_net(d) / 1000) * 24

3. Annual energy (kWh/year):

E_yr = sum_over_days(E_day(d))

4. Location-based emissions (kg CO_2/year):

CO2_yr = sum_over_validators((P_node/1000) * 8760 * EF_c(i))

Where unknown locations use IEA world average EF = 0.459 kg CO_2/kWh.
5. Convert to metric tons:

tCO2_yr = CO2_yr / 1000

Legend:

- P_node: Avg wall power per validator (W), incl. PUE
- EF_c: Grid emission factor (kg CO_2/kWh) for country c
- PUE: Power Usage Effectiveness

4) Testnet Case Study (63 Nodes, Location-Weighted)

Using the 37-node location dataset's average per-node energy and weighted EF, scaled to 63 nodes:

- Per-node annual energy: ~2,306.27 kWh/yr
- Network annual energy: ~145,295 kWh/yr
- Weighted EF: ~0.375 kg CO_2/kWh
- Annual emissions: 54.5 tCO_2/yr

Reconciliation with earlier ~47.5 tCO_2/yr estimate

The earlier ~47 tCO_2/yr figure used a more EU-leaning emission factor (~0.32 kg/kWh) and a slightly different per-node energy assumption (≈2,356 kWh/yr). The current case study applies the location-weighted EF (~0.375 kg/kWh) derived from the 37-node mix and a 2,306 kWh/yr per-node average, then scales to 63 nodes. The methodology now documents these inputs explicitly so changes in EF or per-node energy transparently explain differences in totals.

5) Mainnet Projection (200 Validators)

A) Global-average EF (0.459):

Per-node energy = (269/1000) * 8760 = 2,356.44 kWh/yr
Network energy = 2,356.44 × 200 = 471,288 kWh/yr
CO_2 (t/yr) = (471,288 × 0.459) / 1000 = 216.2 tCO_2/yr

B) Location-weighted EF (~0.375):

CO_2 (t/yr) = (471,288 × 0.375) / 1000 = 176.7 tCO_2/yr

6) CCRI/Solana Methodology Alignment

- **Node Power Measurement:** Commitment to empirical wall power measurement for at least one active and one standby validator, similar to CCRI's multi-configuration approach.

- **Inactive Validators:** Includes powered inactive/candidate validators in N(d) to meet MiCA's comprehensive accounting.
- **Geographic EF:** Matches CCRI in applying location-weighted EF and uses IEA's 0.459 kg CO_2/kWh as fallback.
- **TPS Adjustment:** Added to the mainnet roadmap post-launch, measuring marginal energy per transaction.
- **Hardware Transparency:** Will publish hardware measurement table, replacing modeled values with empirical ones over time.

7) Improving Accuracy

- Daily IP geolocation for EF weighting
- Empirical wall power for representative nodes
- Track standby/inactive node utilization
- Update P_node when hardware specs change
- Integrate TPS vs. power regression post-mainnet

Summary:

- 63-node testnet (location-weighted): ~54.5 tCO_2/yr
- Projected 200-validator mainnet: ~176–216 tCO_2/yr
- MiCA-aligned, CCRI-informed, and ready for live monitoring integration

Year one estimates:

0G Validator Layer — Year-1 Scope 2 (MiCA-aligned)

Inputs

- Validators: 121 (all running)
- Per-node wall power: 269 W → 2,356.44 kWh/year per node
- Method: location-based EF where known; unknowns set to non-EU EF = 0.55 kg CO_2/kWh (conservative)

Location mix used (64 known from your list)
Europe (incl. UK/UA) 40 • North America 9 • Japan 4 • Korea 4 • China 2 • Singapore 2 • Turkey 2 • Australia 1
(57 validators unknown → treated as non-EU at EF 0.55)

Emission factors (kg CO_2/kWh)
EU 0.30 • NA 0.40 • JP 0.46 • KR 0.50 • CN 0.65 • SG 0.52 • TR 0.47 • AU 0.70 • Unknown (non-EU) 0.55

Results

- Total energy (121 nodes): 121 × 2,356.44 = 285,129.24 kWh/year
- Known 64 validators: 55.2 tCO_2/year
- Unknown 57 validators (EF 0.55): 73.9 tCO_2/year
- Total (conservative): 129.1 tCO_2/year

Sense-check / fallback
Using global-average EF 0.459 for all 121 nodes would yield 130.9 tCO_2/year. This approach is similarly conservative and avoids assigning unknowns to low-EF EU grids.

0G Storage Layer Emissions Methodology – Testnet Case Study & Mainnet-Ready Approach

1. Scope and Purpose

This methodology estimates Scope 2 (electricity-related) emissions from the 0G network's storage layer, using the current testnet as a case study. It is designed to be defensible under the MiCA framework and scalable to mainnet.

The method outputs:

- Total network emissions (tCO_2/year or reporting period)
- Emissions intensity per terabyte ($kgCO_2$/TB·year)

It accounts for:

- Active miner counts (from StorageScan)
- Total storage size (TB)
- Power consumption per miner (including facility overhead/PUE)
- Uptime
- Grid carbon intensity (location-based where possible)
- Proof of Random Access (PoRA) overhead

2. Data Sources

Primary: https://storagescan-galileo.0g.ai/

- Active miners (M)
- Total storage size (T) in GB/TB

Secondary:

- Miner location data (when available) via IP geolocation or attestation
- Grid carbon intensity factors from IEA/ElectricityMaps
- Measured miner power profiles (preferred) or industry-standard estimates

3. Methodology

3.1 Boundaries

- **Scope:** Storage miners/providers only (excludes validators, DA, compute layers)
- **Emissions category:** Scope 2 (electricity)
- **Uptime:** 24/7 (8760 hours/year)
- **Power per miner:** Baseline 200 W (wall draw, includes facility PUE). Replace with measured averages as available.
- **Grid factor:** 0.4 $kgCO_2/kWh$ (global average). Replace with country-specific factors when miner locations are known.
- **PoRA overhead:** Included in baseline; apply separate uplift if measured.

Legend of Abbreviations

- **M(d):** Number of active miners on day d
- **T(d):** Total logical storage in TB on day d
- **P_net(d):** Total network power draw in watts on day d
- **P_node:** Average power draw per miner in watts (including PUE)
- **E(d):** Total network energy usage in kWh per day
- **grid_factor:** Carbon intensity factor of the electricity grid ($kgCO_2/kWh$)
- **W_per_TB:** Power draw per terabyte of storage (watts/TB)

- **kWh_per_TB_year:** Annual energy usage per terabyte of storage (kWh/TB/year)
- **CO2_per_TB_year:** Annual carbon emissions per terabyte of storage ($kgCO_2$/TB/year)

3.2 Formulas

For each day d:

1. **Total network power (W):**

- P_net(d) = M(d) * P_node

Calculates the total instantaneous electrical power draw of the entire storage network based on the number of active miners and their average power consumption.

2. **Energy/day (kWh):**

- E(d) = (P_net(d) / 1000) * 24

Converts total network power from watts to kilowatts and multiplies by hours in a day to get total daily energy usage.

3. **Emissions/day ($kgCO_2$):**

- CO2(d) = E(d) * grid_factor

Applies the appropriate carbon intensity factor to daily energy use to estimate daily emissions.

Per-TB intensity:

4. **W/TB:**

- W_per_TB = P_net(d) / T(d)

Calculates the power draw per terabyte of logical storage in the network.

5. **kWh/TB·yr:**

- kWh_per_TB_year = (W_per_TB * 8760) / 1000

Annualizes the per-TB power draw into yearly energy consumption per terabyte.

6. **$kgCO_2$/TB·yr:**

- CO2_per_TB_year = kWh_per_TB_year * grid_factor

Converts annual per-TB energy usage into annual per-TB carbon emissions.

3.3 PoRA Increment

Two options:

1. **Measured:** Instrument reference miner, measure idle vs. PoRA-active over 48–72 h.
2. **Interim:** Apply conservative +10% uplift to P_node network-wide.

3.4 Mainnet Procedure (Daily → Monthly)

1. Daily snapshot of M(d) and T(d) from StorageScan API.
2. Compute daily kWh, $kgCO_2$, and per-TB intensities.
3. If location data is available, apply location-specific grid factors and sum for network total.
4. Aggregate to monthly/annual reports.

4. Testnet Case Study (Using Average Miner & Storage Values)

Inputs:

- Average active miners (M) = 500

- Total storage size (T) = 2.26 TB
- Power per node (P_node) = 200 W (includes PUE)
- Grid factor = 0.4 kg/kWh
- Testnet runtime = ~4 months (~0.333 year)

Network-level annualized totals:

- P_net = 500 * 200 = 100,000 W = 100 kW
- Annual_energy = 100 * 8760 = 876,000 kWh/year
- Annual_emissions = 876,000 * 0.4 = 350,400 kg = 350.4 tCO_2/year

Per-TB intensity:

- W_per_TB = 100,000 / 2.26 ≈ 44,247.8 W/TB
- kWh_per_TB_year = (44,247.8 * 8760) / 1000 ≈ 387,588.5 kWh/TB·yr
- CO2_per_TB_year = 387,588.5 * 0.4 ≈ 155,035.4 kg = 155.0 tCO_2/TB·yr

Adjusted for Testnet runtime:

- Total_emissions_to_date = 350.4 * 0.333 ≈ 116.8 tCO_2
- Per_TB_emissions_to_date = 155.0 * 0.333 ≈ 51.6 tCO_2/TB

5. Observations & Regulatory Considerations

- Per-TB intensity is high due to low utilization (2.26 TB across 500 always-on miners).
- MiCA-compliant reporting can present both absolute emissions and intensity metrics.
- Accuracy improves as:
  - Miner location data becomes available (location-based factors)

- Measured power profiles replace estimates
  - PoRA overhead is quantified
  - Daily snapshots replace static averages
- For mainnet, this methodology can run automatically with StorageScan API feeds.

6. Accuracy Improvement Options

- Miner self-reports: location, hosting type, renewable sourcing, uptime
- Direct measurement: 48–72 h wall power logging for representative miners
- PoRA measurement: quantify challenge-induced load
- Redundancy factor: clarify physical vs logical TB for reporting both metrics
- Operate our own storage miner to:
  - Capture real-world power usage data
  - Verify PoRA overhead in practice
  - Provide baseline for scaling assumptions
- Collect peer location data to improve grid factor accuracy

0G Compute Layer – Scope 2 Emissions Methodology (MiCA-Aligned)

1) Scope & Purpose

This methodology defines a **MiCA-aligned framework** for estimating Scope 2 emissions (electricity) for the **compute layer** of 0G, which provides decentralized AI inference and training workloads.

It is designed to be:

- Repeatable and regulator-ready
- Capable of evolving from **assumption-based models** to **real-time measurement**
- Integrated into the ecoBridge sustainability infrastructure layer

Boundary: All active compute nodes in the 0G compute layer, regardless of whether they are running inference, training, or idle but online.
Excludes: Data Availability (DA) nodes, alignment nodes, validator layer, and storage layer.

2) Compute Layer Context from 0G Docs

From 0G documentation:

- The compute layer executes **decentralized AI inference and compute** with cryptographic verifiability
- Runs in a **trustless execution environment**
- Supports both **model inference** (short-duration workloads) and **model training** (long-duration, high-load workloads)
- **Scaling** is dynamic based on demand

The architecture implies heterogeneous workloads and hardware, likely GPU-based for high-performance AI tasks.

3) Base Calculation Methodology

Formula 1 – Total Network Power (W)

$$P_{net}(d) = \Sigma_i \left(P_{node}(i) \times U(i)\right)$$

Where:

- **$P_{node}(i)$:** Wall power draw of compute node *i* (W), incl. PUE
- **$U(i)$:** Utilization factor for node *i* (0–1 scale: idle ≈ 0.1–0.2, inference ≈ 0.5–0.7, training = 1.0)

Explanation:
This calculates the total instantaneous power demand (in watts) of the entire compute network on a given day **d**.

- It sums the power of each compute node **i**, adjusted by its **utilization factor (U)**, which represents how much of its capacity is being used (e.g., idle = low, training = high).
- **$P_{node}(i)$** includes the **Power Usage Effectiveness (PUE)** so it reflects total facility wall power, not just IT load.

Formula 2 – Daily Energy (kWh)

$$E_{day}(d) = \left(P_{net}(d) / 1000\right) \times 24$$

Explanation:
Converts total power from watts to kilowatts (divide by 1000), then multiplies by 24 hours to get total energy consumption in **kilowatt-hours** for one day across the compute layer.

Formula 3 – Annual Energy (kWh/year)

$$E_{yr} = \Sigma_{over\_days} \left(E_{day}(d)\right)$$

If daily breakdown unavailable:

$E\_yr \approx (\Sigma\_i\ (P\_node(i) \times U(i)) / 1000) \times 8760$

Explanation:
Calculates total energy use in a year.

- If you have daily data, sum each day's **E_day(d)** over the year.
- If you only have averages, multiply average total network power by 8,760 hours (hours in a year) to estimate annual energy.

Formula 4 – Location-Based Emissions (kg CO_2/year)

$CO2\_yr = \Sigma\_i\ ((P\_node(i) \times U(i) / 1000) \times 8760 \times EF\_c(i))$

Where:

- **EF_c(i):** Country-specific grid emission factor (kg CO_2/kWh)

Explanation:
This converts annual energy use into carbon emissions in **kilograms** by applying a **country-specific grid emission factor (EF_c)** for each node's location.

- If a node is in a country with cleaner electricity, EF is lower.
- If only global average data is available, the same EF is applied to all nodes.

Formula 5 – Convert to Metric Tons

$tCO2\_yr = CO2\_yr / 1000$

Explanation:
Converts the carbon emissions from **kilograms** to **metric tons** (tCO_2/year), which is the standard reporting unit for MiCA and most ESG frameworks.

4) Initial Assumptions (Pre-Mainnet)

- **Node type split:** 70% inference, 30% training (adjustable once data is available)

- **Power draw estimates (IT load):**
    - Inference node: ~400 W IT (560 W wall with PUE 1.4)
    - Training node: ~1,200 W IT (1,680 W wall with PUE 1.4)
- **Utilization factors:**
    - Inference: 0.6 average load
    - Training: 0.9 average load
- **Grid emission factor (EF):** Global average 0.4 kg CO_2/kWh unless location data is collected
- **Number of compute nodes (N):** Unknown — will be measured at mainnet launch

5) Transparency & Accuracy Improvements

To move from assumptions to precise, MiCA-compliant reporting:

1. **Run a compute node ourselves**
    - Measure real power draw for idle, inference, and training loads
    - Capture workload-based utilization factors
2. **Geolocation data collection**
    - Check whether peer IP discovery is possible in the compute network
    - If not, deploy a **lightweight telemetry client** for voluntary operator installation
        - Captures: country, hardware specs, uptime, utilization, measured wall power
3. **Daily node count snapshots**
    - Track active compute nodes

    - Breakdown by workload type if available

4. **Integration into ecoBridge dashboard**
    - Real-time ingestion of node count, location, utilization, and power draw
    - Automatic location-weighted EF application

6) Data Needed for Best Results

- **Node count (daily)**
- **Hardware type per node** (GPU model, CPU, RAM)
- **Workload classification** (inference vs. training vs. idle)
- **Geographic location** (country-level minimum)
- **Measured wall power draw** for representative hardware
- **Utilization factor** per workload type

7) Roadmap for Implementation

Pre-Mainnet

- Establish baseline power models via owned compute node
- Define default utilization factors per workload type
- Develop telemetry client for voluntary data sharing

Mainnet Launch

- Begin daily snapshots of compute node counts and geolocations (if available)

- Apply power model + EF weighting for preliminary results

Post-Mainnet (Phase 2)

- Integrate live utilization and location data from telemetry or network API
- Update power model with measured values from representative nodes

Full Automation (Phase 3)

- Continuous ingestion of all key data into ecoBridge
- Automated daily MiCA-compliant emissions reporting for compute layer

Baseline-from-One-Node Plan (Compute Layer)

1) What we'll measure (must-haves)

- **Wall power (W)** with an inline meter (e.g., smart plug/kWh logger).
  - Log at **1–5s** cadence; store **W, kWh, timestamp**.
- **Utilization traces**: GPU %, GPU memory GB, CPU %, RAM GB, disk I/O, net I/O.
- **Workload labels** (attach to time ranges):
  - `idle/standby`
  - `inference – light`
  - `inference – heavy`
  - `training – light`

   - `training – heavy`

- **Environment**: ambient temp, PUE assumption (or datacenter PUE if known).

  Output: a **load → power curve** and **kWh/hour** factors per workload class, plus an **idle baseline**.

2) How we'll run it (repeatable protocol)

1. **Hardware**: pick a representative compute rig (e.g., 1× high-end GPU, server-grade CPU, 64–128 GB RAM, NVMe). Record exact model numbers.

2. **Metering**: install a calibrated wall meter; zero the counter.

3. **Workload suite** (minimum 45–60 min each, steady state):

   - Idle Baselining (no jobs, node online)

   - Inference–Light (small model / low QPS)

   - Inference–Heavy (larger model / higher QPS)

   - Training–Light (short fine-tune or small batch)

   - Training–Heavy (sustained high GPU utilization)

4. **Collect**: W, kWh; GPU/CPU/etc. at 1–5s cadence.

5. **Aggregate**: median W, 95% range, and **kWh/hr** per class.

6. **Sanity-check**: repeat one workload to verify variance; note thermals.

3) Turn measurements into per-node factors

For each workload class **c**:

- **Avg power**: P¯c\overline{P}_c (W)

- **Energy rate**: ec=P¯c/1000e_c = \overline{P}_c/1000 (kWh/hour)

- **If PUE applies** and meter is IT-side, multiply by PUE to get **wall**. (If we meter at the wall, PUE is already captured.)

Define a **daily mix** (fractions that sum to 1):
$\pi = \{ \pi_{\text{idle}}, \pi_{inf,light}, \pi_{inf,heavy}, \pi_{train,light}, \pi_{train,heavy} \}$

Then **per-node daily energy**:
$E_{\text{day,node}} = \sum_c \pi_c \cdot e_c \cdot 24$

Per-node annual energy: $E_{\text{yr,node}} = E_{\text{day,node}} \cdot 365$

Per-node emissions (location-based):
$\text{tCO}_2{}_{\text{node}} = (E_{\text{yr,node}} \times EF_{\text{country}})/1000$

4) Extrapolate to the network (until broader telemetry exists)

- Let **N(d)** = active compute nodes on day d (unknown → scenario ranges: low/med/high).
- If we can't observe each node's country yet, use **IEA 0.459 kgCO$_2$/kWh** fallback; otherwise apply **country EF weighting**.
- **Daily network energy**: $E_{\text{day,net}} = N(d) \times E_{\text{day,node}}$
- **Annualized** (for a reporting period T days): $E_{\text{yr,net}} \approx \frac{\sum_{d=1}^{T} N(d)}{T} \times 365 \times E_{\text{day,node}}$

Provide **scenarios** (e.g., N=50/200/500) and **two workload mixes** (utilization-light vs utilization-heavy) to bound results.

5) What will be published

- Hardware bill of materials (models, firmware/driver versions).
- Meter spec & calibration notes.

- Raw CSVs (timestamped W/kWh + utilization).
- Aggregation notebook (reproducible).
- Final tables: P̄c\overline{P}_c, ece_c, workload mix, Eday,nodeE_{\text{day,node}}, Eyr,nodeE_{\text{yr,node}}.
- EF table (fallback + any location weights used).
- Uncertainty bands (± from repeated runs).

6) How this meshes with DA, Storage, Alignment

- **DA**: identical structure; lighter IT loads; we already defined assumed W and U. Baseline by running one DA verifier and measuring as above.
- **Storage**: baseline two SKUs (HDD‑heavy vs NVMe‑heavy) under PoRA challenge cadence + idle; derive per‑TB‑month kWh factor.
- **Alignment**: extremely low‑spec; measure one SBC/VM instance; likely 1–5 W average. Scale by node count.

0G Data Availability (DA) Layer – Scope 2 Emissions Methodology (MiCA-Aligned)

1) Scope & Purpose

This methodology defines a **MiCA-aligned framework** for estimating **Scope 2 emissions** (electricity) for the **Data Availability (DA) layer** of 0G, which ensures short-term data accessibility for L2 rollups and real-time applications.

It is designed to be:

- **Repeatable and regulator-ready**
- **Capable of evolving from assumption-based to measurement-based**
- **Integrated into the ecoBridge sustainability infrastructure layer**

Boundary: All DA nodes — including **verifiers**, **retrievers**, and **clients** — running continuously, regardless of load level.
Excludes: GPU-based DA encoders unless they operate continuously (treated separately if frequent).

2) DA Layer Context (0G Documentation)

From 0G documentation:

- DA nodes **verify, sign, and store blob data** for immediate retrieval.
- Hardware recommendations:
    - **DA Verifier:** 16 GB RAM, 8-core CPU, 1 TB NVMe SSD, 100 Mbps
    - **DA Retriever/Client:** 8 GB RAM, 2-core CPU, 100 Mbps

  - **DA Encoder:** GPU-enabled (tested on RTX 4090)
- Expected **continuous uptime** for reliability.

3) Calculation Method

Let **N_type(d)** = number of online DA nodes of a given type on day **d**.
Let **X** = total number of nodes of each type (to be determined once operational data is available).

Formula 1 – Total Network Power (W):

$$P\_net(d) = \Sigma\_type \; (N\_type(d) \times P\_node(type) \times U\_type)$$

Explanation:
Sums the total power draw for all DA nodes online on day **d**, adjusting for utilization level (**U_type**) and average wall power (**P_node(type)**).

Formula 2 – Daily Energy (kWh):

$$E\_day(d) = (P\_net(d) / 1000) \times 24$$

Explanation:
Converts total network power from watts to kilowatts, then multiplies by 24 hours to get daily energy.

Formula 3 – Annual Energy (kWh/year):

$$E\_yr = \Sigma\_over\_days(E\_day(d))$$

If daily breakdown unavailable:

$$E\_yr \approx (\Sigma\_type \; (N\_type \times P\_node(type) \times U\_type) / 1000) \times 8760$$

Explanation:
Adds up daily energy across all days in a year. If daily data is unavailable, uses average counts and utilization to estimate annual consumption.

Formula 4 – Location-Based Emissions (kg CO_2/year):

$$CO2\_yr = \Sigma\_nodes\ (\ (P\_node(type) \times U\_type / 1000) \times 8760 \times EF\_c(i)\ )$$

Explanation:
Multiplies each node's annual energy by the **grid emission factor** (**EF_c(i)**) for its country to get emissions in kilograms of CO_2.

Formula 5 – Convert to Metric Tons:

$$tCO2\_yr = CO2\_yr / 1000$$

Explanation:
Converts kilograms of CO_2 to metric tons for reporting.

4) Initial Assumptions (Pre-Mainnet)

Node Type	IT Load (W)	Wall Power (W) w/ PUE=1.4	Utilization Factor (U)	Notes
DA Verifier	36 W	50 W	0.7	8-core CPU, moderate load
DA Retriever	14 W	20 W	0.5	2-core CPU, lightweight tasks
DA Client	14 W	20 W	0.5	Similar to retriever
DA Encoder*	~300 W GPU	420 W	Variable	Only if frequent/continuous

Emission Factor (EF): 0.5 kg CO_2/kWh (global average) unless node location known.

5) Example Annualized Calculation (Hypothetical)

Assume:

- **X verifiers** @ 50 W, U=0.7
- **Y retrievers** @ 20 W, U=0.5
- **Z clients** @ 20 W, U=0.5
- EF = 0.5 kg CO_2/kWh

1. **Verifier total power:** `X × 50 × 0.7` W
2. **Retriever total power:** `Y × 20 × 0.5` W
3. **Client total power:** `Z × 20 × 0.5` W
4. **Total network power:** Sum of above
5. **Daily energy:** `(P_net / 1000) × 24` kWh/day
6. **Annual energy:** `E_day × 365` kWh/year
7. **Emissions:** `(Annual Energy × EF) / 1000` tCO_2/year

6) Roadmap for Accuracy Improvement

Pre-Mainnet:

- Define baseline power models from reference hardware
- Prepare telemetry client for voluntary operator installation

Mainnet Launch:

- Daily snapshots of node counts by type
- IP geolocation for EF weighting

Phase 2 (Post-Mainnet):

- Collect real wall power measurements from representative nodes
- Integrate utilization data (idle vs. peak)

Phase 3 (Full Automation):

- Continuous ingestion of live node, utilization, and EF data
- Automated MiCA-compliant daily reporting

0G Alignment Nodes – Scope 2 Emissions Methodology (MiCA-Aligned)

1) Scope & Purpose

This methodology defines a **repeatable, regulator-ready** framework for estimating **Scope 2 (electricity)** emissions from **Alignment Nodes** in the 0G network.

- **Boundary:** All Alignment Nodes that are online (always-on watchers of network/model behavior).
- **Out of scope:** Validators, Storage, DA, Compute (handled in their own methods).

Alignment nodes are designed to be **lightweight** (laptop/desktop/phone/VM-class), with minimal CPU/RAM and no GPU requirement. Per-node energy is expected to be **very low** compared to validator or compute nodes. The aggregate footprint depends primarily on **how many** alignment nodes are online.

2) Alignment Layer Context

- Primary role: **monitor** protocol and model behavior (e.g., compliance, anomalies, drift), flagging issues.
- Typical hardware: **very low-spec**; no dedicated GPU; small RAM/CPU; small disk.
- Expected uptime: **24/7** (to be ready to attest/flag events).
- Emissions drivers: **node count** × **tiny average power** × **hours** × **grid EF**.

3) Calculation Method

Let **N(d)** = number of online alignment nodes on day **d**. Let **P_node** be the **average wall power** per node (W), already including any PUE (if applicable). If you later track utilization, keep a utilization factor **U** (default **U = 1.0** when P_node is already an average).

Formula 1 – Total Network Power (W)

$P_{net}(d) = N(d) \times P_{node} \times U$

Explanation: Instantaneous electrical power drawn by all alignment nodes online on day d. If P_node is an average (not peak), set U=1.

Formula 2 – Daily Energy (kWh)

$$E\_day(d) = (P\_net(d) / 1000) \times 24$$

Explanation: Convert watts to kilowatts and multiply by 24 hours to get daily energy.

Formula 3 – Annual Energy (kWh/year)

$$E\_yr = \Sigma\_over\_days [E\_day(d)]$$

If daily snapshots are unavailable:

$$E\_yr \approx (N\_avg \times P\_node \times U / 1000) \times 8760$$

Explanation: Sum daily energy across the year, or estimate with an annual hours multiplier when using average counts.

Formula 4 – Location-Based Emissions (kg CO_2/year)

$$CO2\_yr = \Sigma\_nodes [(P\_node \times U / 1000) \times 8760 \times EF\_c(i)]$$

Explanation: Multiply each node's annual kWh by its country grid **emission factor** EF_c(i). If locations are unknown, apply a fallback EF.

Formula 5 – Convert to Metric Tons

$$tCO2\_yr = CO2\_yr / 1000$$

Explanation: Convert kg CO_2 to tCO_2 for reporting.

Legend

- **N(d)** = online alignment nodes on day d
- **N_avg** = average online alignment nodes over period

- **P_node** = average wall power per node (W), inclusive of PUE
- **U** = utilization factor (1.0 if P_node is an average)
- **EF_c(i)** = grid emission factor (kg CO_2/kWh) for node i's country

4) Initial Assumptions (Pre-Mainnet)

Because alignment nodes are intentionally minimal:

Parameter	Baseline assumption	Notes
Average wall power per node (**P_node**)	**5 W**	Conservative; reflects tiny CPU/RAM, no GPU
Utilization (**U**)	**1.0**	P_node is already an average
Uptime	**24/7** (8760 h/yr)	Adjust if measured duty cycle differs
Emission factor (fallback)	**0.459 kg CO_2/kWh**	IEA world-average fallback (CCRI-consistent); swap to location-weighted when IP/location known

Sensitivity: if reality is **1–3.5–8–10 W** per node, totals scale **linearly**.

5) Case Study – Maximum 175,000 Alignment Nodes (Annual)

Inputs

- **N** = 175,000 nodes
- **P_node** = 5 W (= 0.005 kW)
- **U** = 1.0 (average already)
- **Hours** = 8760

- **EF (fallback)** = 0.459 kg CO_2/kWh

Step‑by‑step

P_net = 175,000 × 5 W × 1.0 = 875,000 W = 875 kW
E_day = 875 kW × 24 = 21,000 kWh/day
E_yr = 21,000 × 365 = 7,665,000 kWh/year
CO2_yr (kg) = 7,665,000 × 0.459 = 3,518,235 kg
tCO2_yr = 3,518,235 / 1000 = 3,518.2 tCO_2/year

Result (Base): ~3,518 tCO_2/year (fallback EF 0.459).
Sensitivity (same EF):

Avg W per node	E_yr (GWh)	tCO_2/yr
1 W	1.533	**704 t**
3.5 W	5.366	**2,462.8 t**
5 W	7.665	**3,518.2 t**
8 W	12.264	**5,629.2 t**
10 W	15.330	**7,036.5 t**

Once we geolocate nodes (country‑level), replace the fallback EF with **location‑weighted EF** to tighten accuracy (often reduces tCO_2 if many nodes sit on cleaner grids).

6) Roadmap for Accuracy Improvement

Pre‑Mainnet

- Define the **reference hardware** profile(s) for alignment nodes.
- Measure **empirical wall power** on at least one representative node (smart plug or DC power meter).

- Prepare a **lightweight telemetry** option (opt‑in) to report: country, uptime, wall power (if measured), OS‑level CPU/network utilization.

Mainnet Launch

- Take **daily snapshots** of active alignment node counts (N(d)).
- Attempt **IP geolocation** (country‑level) to apply location‑based EF.
- Publish a **baseline power table** (measured vs. assumed).

Post‑Launch (Phase 2)

- Ingest **telemetry samples** to refine P_node (mean/median, percentiles).
- If utilization materially varies, reintroduce **U** as a dynamic factor (idle vs. spike).
- Update EF weighting with new location coverage.

Full Automation (Phase 3)

- Continuous ingestion of **N(d), location, and P_node** (measured or imputed).
- **Daily** MiCA‑ready Scope 2 reporting for alignment nodes, with uncertainty bands.