Supplementary Note S1—Scenario construction for Figure 1

Purpose and scope

Figure 1 presents a transparent and reproducible envelope of the global electricity demand attributable to data centres for the period 2022–2035. Starting from a small set of totals anchored in literature and from presented assumptions, it was decided to construct three internally consistent scenarios (Low, Central, High) together with a shaded envelope that bounds their evolution. The purpose is threefold:

(i) to anchor the discussion of the Artificial Intelligence (AI) energy paradox in the quantities used by system operators and planners (TWh, GW),

(ii) to test policy proposals against plausible trajectories that can be followed by such systems, and

(iii) to disclose all the moving parts so the reader may stress–test the assumptions.

Summarizing, the figure serves as a compact reference frame for the problem in our days.

Inputs and sources

- Baseline year total electricity for the combined category:

$$E_{2022} = 460 \text{ TWh}$$

- 2026 calibration band:

$$E_{\{2026\}} \in [620, 1050] \text{ TWh with a central value of 800 TWh}$$

Source: International Energy Agency (IEA), Electricity 2024; combined series "Global electricity demand from data centres, AI and cryptocurrencies".

All other quantities listed below are assumptions.

Assumptions

- PUE control points (2022, 2026, 2030, 2035) for each scenario are linearly interpolated between points.
- Hardware efficiency gains e_s (%/y): Low 15, Central 18, High 22 (defaults; tunable in excel sheet).
- Utilization U_s (duty-factor proxy): Low 0.45, Central 0.36, High 0.30 (defaults; tunable in excel sheet).

Outputs

- Lines in the graph: E_t for Low/Central/High (TWh).

Methodological Approach

Annual Facility Electricity

The Annual Facility Electricity (E_t) is modeled as:

$$E_t = PUE_t \cdot E_0^{IT} \cdot \left(\frac{1+g}{1+e}\right)^t$$

with:

$$E_0^{IT} = \frac{E_{\{2022\}}}{PUE_{\{2022\}}}$$

Where:

PUE: Power Usage Effectiveness

g: workload growth (compute demand)

Calibration to 2026

For each scenario $s \in \{Low, Central, High\}$ we back-solve the workload growth g_s required to hit the 2026 anchor and PUE_t is back-solved:

$$r_s = \left(\frac{E_{2026}^{(S)}}{E_{2022} \cdot \frac{PUE_{2026}^{(s)}}{PUE_{2022}^{(s)}}}\right)^{\frac{1}{2026-2022}}$$

$$g_s = r_s(1 + e_s) - 1$$

Where:

r_s: net energy growth

e_s: hardware efficiency gain (performance/W)

This ensures $E_{2026}^{(s)}$equals the literature value by construction.

Facility power

The average power is:

$$\overline{\{P\}}_t = \frac{E_t}{8760}\text{TWh}$$

For planning-type capacity comparisons an indicative nameplate is also calculated

$$P_t^{\{name\}} = \frac{\overline{\{P\}}_t}{U_s}$$

Using a constant scenario-level utilization U_s (a duty-factor proxy derived from training/inference shares and cycle factors).

- Shaded envelope: $E_t^{High} - E_t^{Low}$
- Markers: literature anchors at 2022 and 2026 (central, plus low/high at 2026).
- Derived series in SI: $\overline{\{P\}}_t$(GW) and $P_t^{\{name\}}$(GW) for each scenario
- Excel file: Figure1.xlsx

Limitations

The envelope is scenario-based (not probabilistic). It abstracts from price–demand feedbacks, geography, and network constraints.

Excel Sheet



Figure1.xlsx