

# How Much Electricity Does Bot Web Traffic Actually Use?

A transparent order-of-magnitude bound - v0 (proof-of-method, not precision)

Author: N. Milton · ORCID: [0009-0003-4213-7769](https://orcid.org/0009-0003-4213-7769) Version: v0 · Date: 2026-06-02 · License: CC BY 4.0

## Abstract

Public claims about the electricity consumed by automated ("bot") web traffic vary by more than an order of magnitude, yet are routinely presented as single figures. This paper shows that no single honest figure is recoverable from public data, and that the headline numbers in circulation (on the order of ~250 TWh/yr) are inflated roughly 15-20x by three implicit methodological choices: (1) multiplying a request-share by *total* data-centre energy, a pool now dominated by non-request workloads such as AI training; (2) using a web-application bot fraction (~53%) in place of a whole-network one (~31%); and (3) allocating energy by *average* share rather than the *marginal* energy actually saved on a baseload system. Carving the relevant energy pool, using a frame-explicit bot-fraction range, and applying marginal allocation yields a transparent bound of **~5-50 TWh/yr** (central estimate **~15 TWh**, ~0.05% of global electricity). Two independent cross-checks - a bottom-up request-level estimate and a non-circular comparison with bot-mitigation market spend - agree within one order of magnitude. The contribution is methodological: a reproducible "deflation" procedure built entirely from free, primary sources (Cloudflare Radar, IEA), plus an explicit account of which assumptions move the result most. All inputs, the computation, and the live data pull are included for reproduction.

**Keywords:** bot traffic; data-centre energy; AI energy; order-of-magnitude bound; deflation; marginal vs average allocation; Cloudflare Radar; IEA.

## The one-line claim

*The honest annual energy attributable to bot web traffic is **~5-50 TWh** (best-estimate centre **~15 TWh**), **not the ~250 TWh** you get by multiplying "bots are half of traffic" by total data-centre energy. The 15-20x gap is **entirely methodological** - three definitional choices, each of which the scary version gets wrong.*

For scale: ~15 TWh ~ 0.05% of global electricity (~30,500 TWh). The viral framing inflates this to ~0.8%.

## The bound, built as a waterfall (watch each defensible choice deflate it)

bot energy ~ relevant energy pool x bot fraction x allocation factor

Step (each = one defensible methodological choice)	Pool (TWh)	x Bot frac	x Alloc	= TWh/yr
<b>Naive / viral</b> - full DC energy x web-app bot % x average	485	0.53	1.0	<b>~257</b>
<b>+ carve the pool</b> - request-serving + server-side network only (NOT training)	~235	0.53	1.0	~125
<b>+ honest bot frame</b> - Cloudflare whole-HTTP, not Imperva web-app	~235	0.31	1.0	~73
<b>+ marginal allocation</b> - energy <i>actually saved</i> , not average share	~235	0.31	0.2	<b>~15</b>

Step (each = one defensible methodological choice)	Pool (TWh)	x Bot frac	x Alloc	= TWh/yr
<b>Full honest band</b> - all three params across their real ranges	150-320	0.31-0.53	0.1-0.3	<b>~5-50</b>

**Independent bottom-up check:**  $\sim 5 \times 10^{15}$  bot requests/yr x  $\sim 0.5$ -5 J/request  $\sim$  **~1-7 TWh**. Overlaps the top-down band's low end ( $\sim 5$ -7 TWh) - the two methods agree where it counts.

## The three parameters that move everything (the deflation, named)

**1. The energy pool - the denominator trap.** Bot fraction is measured on **HTTP requests**. But total data-centre energy ( $\sim 485$  TWh, 2025; IEA) includes large non-request workloads - AI training, storage, internal compute - that serve no external requests. Multiplying a request-share by that whole pool is invalid. Carve to the **request-serving DC slice + server-side network transmission**:

- *Anchor (IEA):* AI-specific servers were only **~53-76 TWh in 2024** ( $\sim 15$ -18% of DC energy);  $\sim 70$ -110 TWh in 2025. So **traditional workloads still dominate at ~80%** ( $\sim 375$ -415 TWh) - the request-serving narrative is about 2030 *growth*, not today's split.

- Request-serving share of that traditional  $\sim 375$ -415 TWh  $\sim$  **~110-210 TWh**; + server-side network  $\sim$  **~50-120 TWh** -> **pool  $\sim$  ~160-330 TWh**. (Data-transmission networks  $\sim 260$ -360 TWh, IEA 2022, are a *separate* pool;  $\sim ?$  is mobile last-mile radio - carved out.)

- **This carve roughly halves the number - and it matters MORE over time:** AI is  $\sim$  half of DC *growth* to 2030, so a 2030 bound needs a far harder carve. v0 is the forgiving 2025 snapshot; do not extrapolate the pool forward unchanged.

**2. The bot fraction - the sampling-frame trap.** "Bots are 53% of traffic" (Imperva) is a **web-application** sample, where bots concentrate. Across *all* HTTP requests (Cloudflare Radar,  $\sim 81$ M req/s), it's **~31%**. Neither is wrong - they measure different denominators. The true figure is a **frame-dependent range, ~31-53%**, not a point. Use the range; name the frame.

**3. The allocation factor - the headline deflation.** Data centres run on **baseload**. Removing bot traffic does **not** linearly remove energy - you can't power down servers you keep for peak human load. The energy *actually saved* (marginal) is a fraction ( $\sim 0.1$ -0.3x) of the bot's *average* share. Most scary bot-energy numbers quietly assume average allocation (factor = 1.0). **This is a ~3-10x sleight of hand, and making it explicit is this analysis's core contribution.**

## Cross-check: does the energy bound agree with what firms spend? (non-circular)

The energy estimate was built **blind** from compute/network data. Independently, the **bot-mitigation market** (revealed-preference: firms defend up to  $\sim$  the harm avoided) is **~\$0.8-1.3B/yr**.

- Energy bound 5-50 TWh x  $\sim \$70$ /MWh  $\sim$  **\$0.35-3.5B/yr** in electricity.

- Defensive floor  $\sim$  **\$0.8-1.3B/yr**.

**Same order of magnitude.** Two independent rails - physical compute and economic defence spend - land within  $\sim 1$  OOM of each other. That coherence is the strongest evidence the bound is honest. (Had the energy number come out at \$100B, we'd know something was inflated.)

## Why no single honest number exists (and never will, from public data)

Only **hyperscalers and CDNs** hold the joined telemetry - traffic *type* x energy *per workload* - needed

to pin this down, and they don't publish it. Everything public is either a **traffic** measurement (no energy) or an **energy** measurement (no traffic breakdown). Bridging them requires the three choices above, and **each choice legitimately moves the answer by a multiple**. So the deliverable is not a figure - it's a **defensible range plus an explicit account of which assumptions move it most**. Anyone quoting a single number is hiding three choices.

---

## Honesty ledger (what this v0 does *\*not\** do)

- **OOM, not precision.** The pool (~160-330 TWh) is the softest input. It is now *anchored* (IEA AI-vs-traditional split), not guessed - but the request-serving fraction of traditional workloads is still modelled, not measured. The EU data-centre disclosure database (Del. Reg. (EU) 2024/1364) firms up *total* and *European* DC energy but reports facility-level KPIs only (PUE, water, renewables) - **no workload-type breakdown**, so it cannot close this gap. Treat the centre as a midpoint, not a measurement.
  - **Value-neutral on purpose.** This bounds *how much energy*, not *how much is wasted*. The "which traffic earned its electricity" question (deadweight vs telemetry vs scraping vs human-serving) is a **separate v2 layer** - kept out of here so the measurement stays credible.
  - **AI-crawler attribution parked.** AI crawlers are ~22% of bot traffic and the growth driver; quantifying the share specifically externalised by named AI firms is a v2 extension, not a v0 claim.
- 

## Reproducibility

All figures in this paper are reproduced by reproduce.py (included), which runs the waterfall, the full-band bound, and both cross-checks from the documented inputs, and can optionally pull the live bot-fraction from the Cloudflare Radar API. See README.md for the exact commands. Primary inputs are free and public; the bot fraction used here was self-pulled from the Cloudflare Radar API on 2026-06-01.

---

## Sources (graded)

- **Primary / load-bearing:** Cloudflare Radar (bot fraction, free API, ~81M req/s) · IEA *Energy and AI / Electricity 2026* (DC energy) · IEA *Data Centres & Data Transmission Networks* (network pool).
- **Vendor / contrast:** Imperva (Thales) *Bad Bot Report* (web-app bot frame - the high end) · bot-mitigation market TAM reports (Fortune Business Insights et al. - defensive floor, take low/narrow end).

*The disagreement between Cloudflare and Imperva is not a problem to resolve - it IS the finding.*

---

## Authorship and AI-assistance disclosure

**Author:** N. Milton (ORCID: 0009-0003-4213-7769).

This analysis was produced with AI assistance, disclosed here in the interest of transparency. The author set the research question, the methodology, and all judgment calls - the choice and ranges of the three parameters, the build order, the deliberately value-neutral scope, and the interpretation of results - and verified every quantitative claim against the cited primary sources. **Anthropic's Claude (model: Claude Opus 4.8)** assisted with data retrieval, calculation, literature search, and drafting. Responsibility for the content, including any errors, rests with the author.

---

## How to cite

*Milton, N. (2026). How Much Electricity Does Bot Web Traffic Actually Use? A transparent order-of-magnitude bound (v0). Zenodo. [https://doi.org/\[DOI assigned on publication\]](https://doi.org/[DOI assigned on publication])*

---

## License

This work is licensed under a **Creative Commons Attribution 4.0 International License (CC BY 4.0)**. You are free to share and adapt the material for any purpose, including commercially, provided you give appropriate credit. Full terms: <https://creativecommons.org/licenses/by/4.0/>