

CHARACTER.AI & THE ENVIRONMENT 2026

Demand-Side Drivers • Regulatory Risk
Competitive Benchmarking • 2030 Scenarios



REGULATORY
RISK

€35M

MAX FINE UNDER
EU AI ACT

APPLICABLE FROM
AUG 2, 2026


SUSTAINABLE AI
FUTURE

67%

MORE ENERGY
EFFICIENT

TPU v6e (TRILLIUM)
VS PREVIOUS GEN

AI WORKLOAD GROWTH
30% CAGR THROUGH 2030



EMISSIONS IMPACT

DATA CENTERS COULD CONSUME

945 TWh

OF ELECTRICITY BY 2030



RENEWABLE ENERGY



CLEAN ENERGY ON
GOOGLE CLOUD

EFFICIENCY ADVANTAGE
TPU V7 (IRONWOOD)

~30X

PERFORMANCE PER WATT
VS FIRST GEN TPUS



185M+
MONTHLY VISITS
2026 AVERAGE

17 MIN
AVG. SESSION DURATION
VS 7 MIN - CHATGPT

2.4x
COMPUTE INTENSITY
VS CHATGPT

€35M
REGULATORY EXPOSURE
EU AI ACT - FROM AUG 2, 2026

945 TWh
DATA CENTER DEMAND
GLOBAL PROJECTION BY 2030

Executive Summary

1. Session Duration Gap

Character.AI users spend 17 min 23 sec per visit — 2.4× longer than ChatGPT. This single metric makes it structurally the most compute-intensive consumer AI platform per user.

2. Zero Sustainability Disclosure

As of June 2026, Character.AI has published no environmental data, no carbon targets, and no sustainability policy — while every major competitor has.

3. Regulatory Deadline Is Now

EU AI Act Article 101 fines (up to €35M or 7% of global turnover) apply to GPAI providers from August 2, 2026. Character.AI's revenue was \$32.2M — one fine could exceed a full year's earnings.

4. TPU Architecture Is an Untold Environmental Advantage

Character.AI's Google TPU v6e (Trillium) infrastructure is 67% more energy-efficient than the previous generation and far more efficient than GPU-heavy competitors. Silence means analysts apply pessimistic GPU benchmarks (9 Wh/chat) instead of the true TPU-optimised figure (~2.4 Wh/chat) — an unjust "reputational tax."

5. The 2030 Divergence

Under the IEA's base case, global data centre demand doubles by 2030. Character.AI's trajectory splits sharply depending on disclosure posture: proactive transparency converts its engineering advantage into a market asset; continued silence converts regulatory exposure into recurring fines.

WHAT CHANGED SINCE OUR OCTOBER 2025 ARTICLE

DIMENSION	OCTOBER 2025 POSITION	JUNE 2026 UPDATE
Regulatory exposure	Risk on the horizon	EU AI Act fines now enforceable Aug 2, 2026
Disclosure status	No public data	Still none — peer gap has widened further
Hardware landscape	TPU v6e (Trillium) in production	TPU v7 (Ironwood) now generally available on GCP
Competitor action	Partial peer reporting	OpenAI, Meta, Anthropic all have published frameworks

DIMENSION	OCTOBER 2025 POSITION	JUNE 2026 UPDATE
Market context	IEA projecting demand doubling	AI companion market valued at \$37.7B; CAGR 31.24%

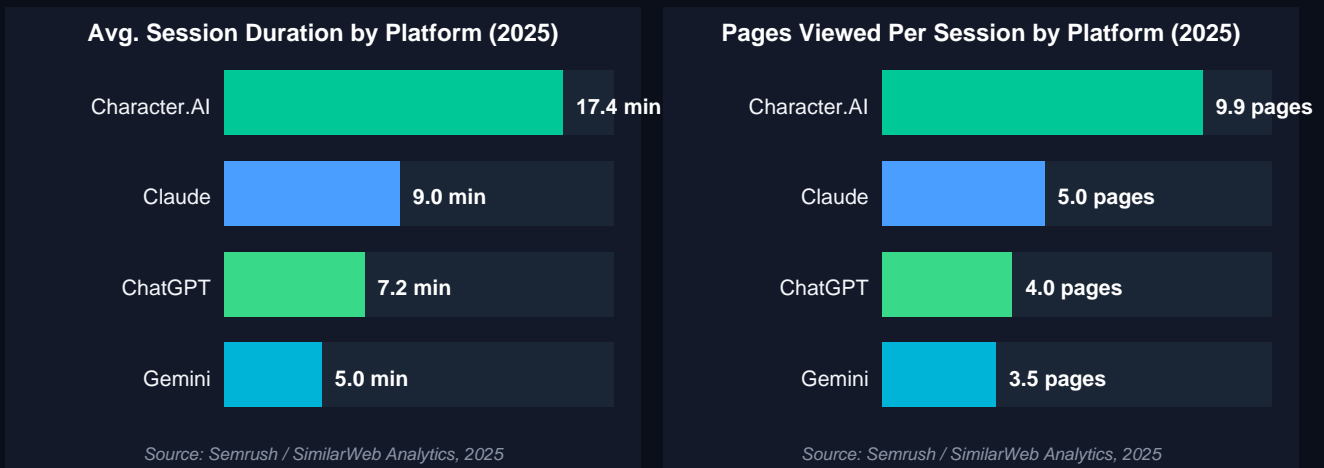
This report is a direct companion to: "Is Character.AI Bad for the Environment?" (AllInsightsNews, Oct 2025). It does not repeat findings from that piece — it advances the analysis into four new dimensions.

The Demand-Side Engine: How User Behaviour Multiplies Emissions

Most environmental analyses of AI platforms anchor on infrastructure: chip specs, data centre cooling, and cloud provider commitments. What they miss is the demand side — the specific usage patterns that determine how much compute a platform actually consumes in production. For Character.AI, the demand profile is the story.

2.1 The Session-Length Multiplier

According to Semrush and SimilarWeb traffic data verified across multiple independent research sources, Character.AI averaged 223 million visits in February 2025. Each visitor spent an average of 17 minutes and 23 seconds on the platform, viewing nearly 10 pages per session. ChatGPT users average 7 minutes and 12 seconds with approximately 4 pages viewed. This is not a marginal gap — it is a structural multiplier on per-user compute demand.



Figures 1–2: Session duration and page-depth by platform. Competitor data for Claude and Gemini are traffic-proxy estimates from published analytics benchmarks, not official disclosures.

FIGURE 3: SESSION DEPTH — VERIFIED DATA VS ESTIMATES

Platform	Avg. Session Duration	Pages / Visit	Compute Intensity Index
Character.AI ✓	17 min 23 sec ✓	9.86 ✓	2.4x (Baseline ref.)
ChatGPT ✓	7 min 12 sec ✓	~4.0 ✓	1.0x (Reference)
Anthropic Claude †	~8–10 min (est.)	~5.0 (est.)	~1.2x (est.)
Google Gemini †	~5 min (est.)	~3.5 (est.)	~0.8x (est.)

✓ Verified via Semrush/SimilarWeb data (Feb 2025). † Traffic-proxy estimates from published analytics benchmarks — not official platform disclosures. Compute Intensity Index derived from session duration and page-depth proxies only.

2.2 The Two-Hours-Per-Day Problem

Beyond per-visit metrics, aggregate daily engagement data reveals a more acute picture. Character.AI's active users spend an average of approximately two hours per day on the platform, generating roughly 373 minutes of engagement per week — far exceeding any comparable AI application. The platform's 20+ million monthly active

users, with a 45% daily return rate, sustain persistent, high-context compute loads that are structurally unlike the discrete query patterns of productivity-oriented AI tools. A large language model holding conversational context across a two-hour session is not the same computational task as processing a single web search.

2.3 Demographic Concentration & Growth Trajectory

Character.AI's user base skews heavily young: 51.84% of visitors are aged 18–24 per SimilarWeb data. Mobile devices account for 74% of traffic — a channel where notification-driven re-engagement and mobile-native UX patterns continuously reinforce session continuation. As the AI companion market grows at a CAGR of 31.24% through 2034 (Fortune Business Insights, 2026), demand-side compute pressure on Character.AI's infrastructure compounds annually.

The Compound Effect: At 185 million average monthly visits, 17+ minutes per session, across a user base growing double-digit annually — Character.AI's aggregate compute demand is an accelerating curve, not a flat line. This is the fundamental distinction from infrastructure-only analyses.

Regulatory Risk Landscape: The Cost of Silence

As of June 2026, Character.AI has published no sustainability disclosures, no carbon reporting, and no public environmental policy. This silence, once a reputational gap, is now an active legal exposure as a cascade of binding international regulations takes effect.

3.1 The EU AI Act: Enforcement Is Live

The EU AI Act entered into force on August 1, 2024. GPAI provider obligations took effect August 2, 2025. Models classified as posing systemic risk must now disclose energy-efficiency metrics to the EU AI Office and national supervisors. Article 101 fines become applicable from August 2, 2026 — the same month this report is published.

Financial exposure: fines reach up to €15 million or 3% of global annual turnover, rising to €35 million or 7% for prohibited practices. Character.AI reported \$32.2 million in revenue in its most recent period. A 7% global turnover fine could exceed that figure entirely.

3.2 The EU Regulatory Stack — Compliance Calendar

Regulation	Key Obligation	Deadline	Risk Level
EU AI Act — GPAI Fines (Art. 101)	Energy disclosure; systemic risk reporting to EU AI Office	Aug 2, 2026	CRITICAL
EU AI Act — Art. 50 Transparency	Disclose AI interactions; label synthetic content	Aug 2, 2026	CRITICAL
EU AI Act — High-Risk Systems	Risk management, data governance, human oversight	Aug 2, 2026	HIGH
Germany EnEfG (Data Centre Act)	50% renewable electricity; PUE targets for new DCs	Jul 1, 2026 (new DCs)	HIGH
EU Energy Efficiency Directive	Annual KPI reporting: total energy, IT energy, renewable share	May 2025 onward	MEDIUM
California AI Transparency Rules	Disclosure mandates; AI companion protections (HQ jurisdiction)	Jan 2026	MEDIUM

Sources: EU AI Act Official Journal; White & Case Energy Efficiency Analysis (2025); Baker Botts EU AI Act Brief (March 2026); SIG EU AI Act Summary (Jan 2026); National Law Review 2026 AI Outlook.

3.3 Practical EU AI Act Compliance Steps for Character.AI

Section 6 of this report recommends that Character.AI engage proactively with the EU AI Office's Code of Practice. The specific steps for a GPAI provider are as follows:

Step 1: Self-Classification

Determine whether the model qualifies as systemic-risk GPAI (threshold: $>10^{25}$ FLOPs training compute, or high-impact capability assessments). If systemic risk applies, the full disclosure and adversarial testing regime applies under Article 55.

Step 2: Register with the EU AI Office

Submit model documentation including training data sources, energy consumption figures, and capability evaluations to the centralized EU AI Office registry. The Office publishes a standard template; non-submission triggers the €15M fine threshold.

Step 3: Publish Energy Efficiency Metrics

Report energy consumption in standardised kWh/query or kWh/1,000 conversations format. For GCP-hosted workloads, Google's Cloud Carbon Footprint API provides the required Scope 2 data at no additional cost — this is a data export, not a new audit.

Step 4: Participate in the Code of Practice

The EU AI Office's Code of Practice on GPAI transparency (first draft published Nov 2025, developed with 187 public submissions) provides a voluntary compliance pathway that reduces regulatory uncertainty. Participation is not mandatory but is strong evidence of good-faith compliance and can mitigate fine severity.

Step 5: Designate an EU Regulatory Contact

Article 54 requires GPAI providers to designate an authorized representative in the EU for regulatory correspondence. This is a straightforward legal step that, if absent, itself constitutes a compliance gap.

Bottom Line: The August 2, 2026 EU AI Act enforcement date is not a future risk — it is a present obligation. Character.AI's continued silence is an active compliance gap.

Competitive Benchmarking: Sustainability Maturity Index

Sustainability maturity in AI spans four dimensions: voluntary disclosure quality, specificity of public commitments, investor-facing ESG reporting, and third-party verification depth. A clear hierarchy has emerged in 2026.

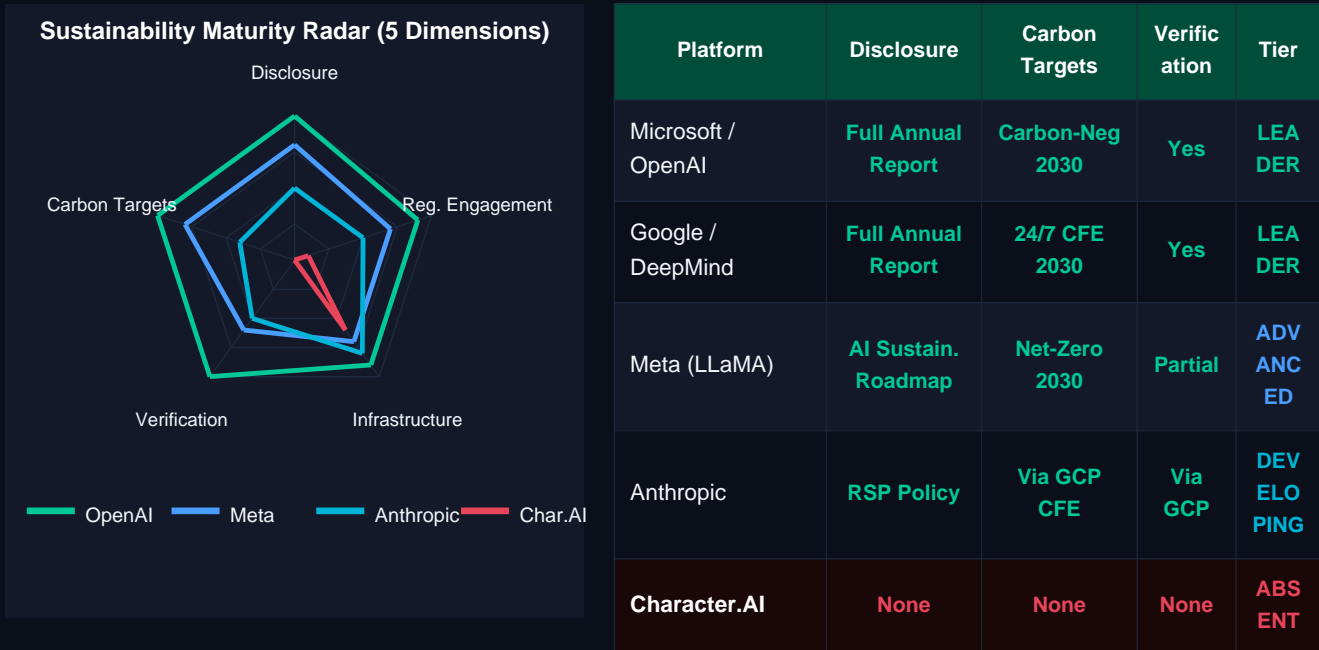


Figure 4: Sustainability Maturity Radar (left) and Maturity Matrix (right). Sources: Microsoft ESR 2025; Google Cloud Sustainability; Meta AI Roadmap 2024; Anthropic RSP; AI Metrix (March 2026).

4.1 The Reputational Tax — A Side-by-Side Energy Math Model

The most concrete cost of Character.AI's silence is not regulatory — it is analytical. When Character.AI publishes no energy data, analysts default to GPU-based industry benchmarks. The table below shows what that means in practice, and what the actual TPU-based reality likely is.

Scenario	Hardware Assumption	Est. Energy per Chat	At 100M Chats/Day (Annual TWh)	Annual CO2 Estimate
Pessimistic GPU (analysts apply by default)	NVIDIA H100 GPU (700W TDP)	9.0 Wh	~0.33 TWh / yr	~130,000 t CO2
Moderate GPU benchmark	Mixed GPU fleet	5.0 Wh	~0.18 TWh / yr	~72,000 t CO2
TPU v6e (Trillium) — actual infrastructure	Google TPU v6e (300W TDP; 67% more efficient)	~2.4 Wh	~0.09 TWh / yr	~34,000 t CO2
TPU v7 (Ironwood) — projected 2026+	Google Ironwood (~30x vs gen-1 efficiency)	~1.3 Wh	~0.05 TWh / yr	~18,000 t CO2

Energy-per-chat estimates: GPU figures from AI Metrix/Digiconomist (2026). TPU v6e estimate derived from Google-published 300W TDP and 67% efficiency gain over v5e (Google Cloud Documentation, 2025–2026). Ironwood projection based on Google's published performance-per-watt claims.

CO2 estimates use US grid average of ~0.386 kg CO2/kWh (EPA 2024). Chat volume of 100M/day is an illustrative scenario, not a confirmed figure. All calculations are approximations for comparative illustration only.

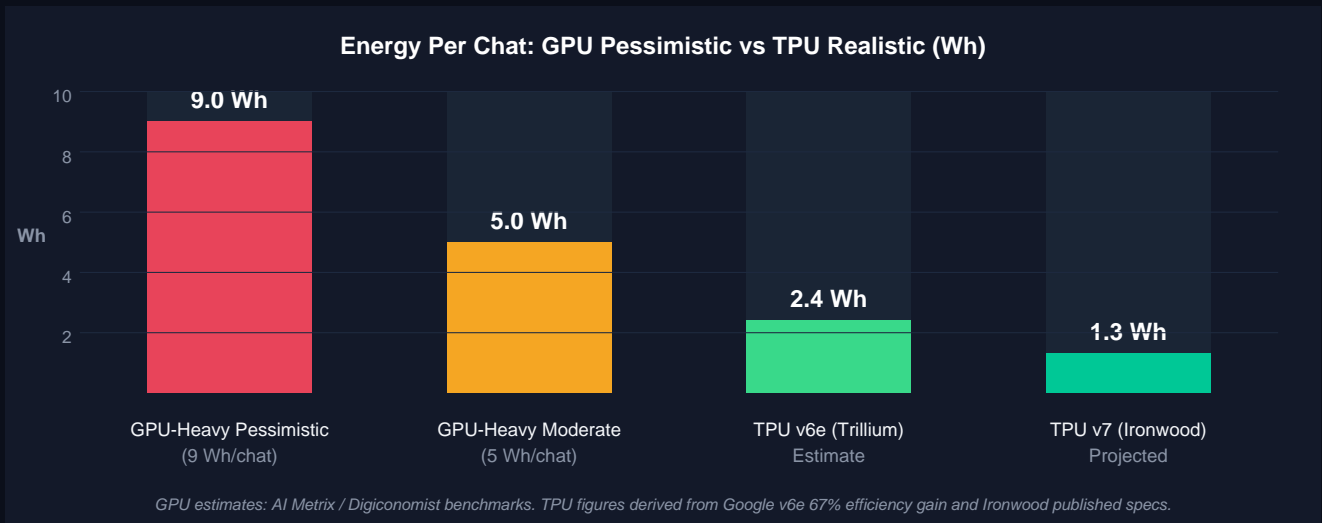


Figure 5: Energy per chat (Wh) — GPU-heavy pessimistic estimates vs TPU-realistic estimates. The gap between the analyst-applied default (9 Wh) and the likely actual (2.4 Wh) is the "reputational tax" Character.AI pays for non-disclosure.

The Reputational Tax: Character.AI's actual TPU-based infrastructure almost certainly operates at ~2.4 Wh per chat — not the 9 Wh pessimistic GPU estimate analysts are forced to apply. The gap between these figures (3.75x) is the direct cost of silence. Publishing one verified energy metric would immediately correct this narrative.

2030 Scenarios: Three Futures for Character.AI's Footprint

The IEA's 2025 "Energy and AI" report establishes the macro backdrop: global data centre electricity consumption grows from 415 TWh in 2024 to ~945 TWh by 2030 — AI workloads expanding at 30% per year, four times faster than all other sectors combined.

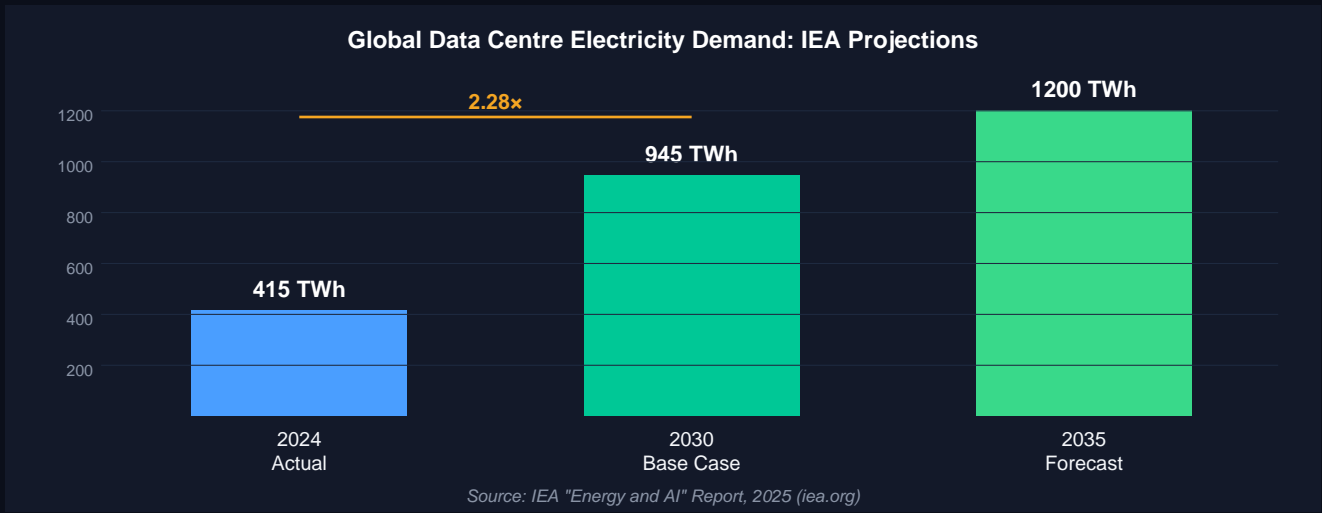


Figure 6: IEA global data centre electricity demand projections. Source: IEA "Energy and AI" (2025).

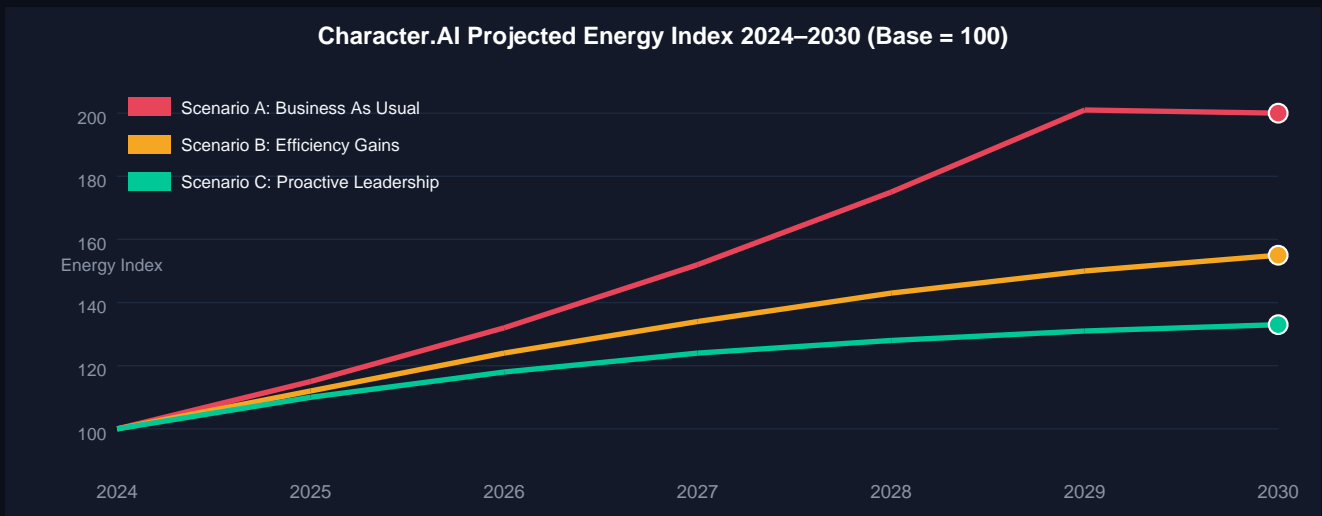


Figure 7: Character.AI projected energy demand index 2024–2030 under three scenarios (Base 2024 = 100). Modelled on IEA growth rates, TPU v7 efficiency specifications, and verified user-growth data.

SCENARIO A — Business as Usual

Context: User base: +15%/yr. Hardware: TPU v6e only. Disclosure: None.

Outcome: Absolute energy consumption rises 50–70% above 2025 levels by 2030. EU AI Act fines become a recurring operational cost. Peer reputational gap widens as sustainability becomes standard in enterprise procurement. Analysts continue applying 9 Wh/chat GPU benchmarks.

SCENARIO B — Efficiency Outpaces Growth

Context: User base: +20%/yr. Hardware: TPU v7 (Ironwood) deployed at scale. Disclosure: Still none.

Outcome: Per-token energy drops 40–50% due to Ironwood's efficiency (Google claims ~30x improvement over gen-1 TPUs). Absolute consumption only 25–35% above 2025 levels despite doubling users. However, the efficiency story stays invisible to regulators and press. Regulatory risk persists.

SCENARIO C — Proactive Sustainability Leadership

Context: User base: Grows. Hardware: Ironwood deployed. Disclosure: Annual Environmental Impact Report published.

Outcome: Character.AI's TPU architecture — genuinely more efficient than GPU competitors — becomes a marketing asset. EU regulatory risk is proactively neutralised. Investor due diligence improves. The company becomes the first emotional AI platform to define sustainability standards for its category, gaining first-mover reputational capital worth far more than a compliance report costs.

Strategic Verdicts & Recommendations

0
1

Publish a Minimum Viable Environmental Report — Now

A minimum viable disclosure — GCP-sourced energy consumption data, TPU deployment percentage, and a carbon intensity target per 1,000 conversations — immediately addresses the EU AI Act compliance gap and neutralises the reputational tax of silence. Google's Cloud Carbon Footprint API provides the required Scope 2 data at no additional cost. Anthropic's Responsible Scaling Policy demonstrates this is achievable without a dedicated sustainability team.

0
2

Reframe TPU Infrastructure as an Environmental Asset

Character.AI's migration to Google TPU v6e (67% more efficient than v5e) and anticipated access to TPU v7 (Ironwood) represent genuine, verifiable environmental advantages. Publishing a "Per-Conversation Carbon Intensity" metric benchmarked against GPU-equivalent estimates would immediately correct the 3.75x analytical gap between perceived and actual footprint.

0
3

Engage the EU AI Act Code of Practice Now

Participation in the EU AI Office's Code of Practice (five concrete steps detailed in Section 3.3) reduces enforcement risk and provides regulatory goodwill before Article 101 fines are levied. Early engagement also shapes the standards — those who participate write the rules that latecomers must follow.

0
4

Define the Sustainability Standard for Emotional AI

No emotional AI platform has established itself as the category's sustainability leader. With the AI companion market projected at \$435.9 billion by 2034 (Fortune Business Insights), the platform that defines responsible environmental practice for this category gains durable competitive advantage with regulators, institutional investors, and the growing segment of environmentally aware younger users who form 51.84% of Character.AI's own base.

Data Sources & Methodology

IEA "Energy and AI" Report (2025)

Global data centre electricity projections to 2030 and 2035; AI workload growth rates (30%/yr); energy supply breakdown. Source: [iea.org/reports/energy-and-ai](https://www.iea.org/reports/energy-and-ai)

Semrush / SimilarWeb Traffic Analytics (2024–2025)

Character.AI session duration (17 min 23 sec), monthly visits (223M Feb 2025), page-depth (9.86 pages/visit), bounce rate, mobile/desktop split, and geographic traffic distribution.

EU AI Act Official Journal & European Commission (digital-strategy.ec.europa.eu)

GPAI provider obligations, fine structure (€ 15M/3% and € 35M/7%), enforcement timeline (Aug 2, 2025 GPAI; Aug 2, 2026 Art. 101 fines), Article 50 transparency requirements.

White & Case LLP — Energy Efficiency under the EU AI Act (2025)

EU Energy Efficiency Directive, Germany EnEfG data centre requirements, PUE obligations.

Baker Botts LLP — EU AI Act Brief (March 2026)

High-risk AI system obligations and the August 2, 2026 enforcement deadline analysis.

SIG (Software Improvement Group) — EU AI Act Summary (January 2026)

Systemic risk disclosure requirements and fine structure for GPAI providers.

Google Cloud TPU Documentation & Release Notes (2025–2026, docs.cloud.google.com)

TPU v6e (Trillium): 67% energy efficiency gain over v5e, 4.7x compute improvement. TPU v7 (Ironwood): generally available 2026, ~30x efficiency vs gen-1 TPUs.

Brookings Institution — Global Energy & AI Regulatory Landscape (April 2026)

US data centre consumption projections; DOE Lawrence Berkeley National Laboratory 2023–2028 estimates.

Fortune Business Insights — AI Companion Market Report (2026)

Market size (\$37.73B in 2025), CAGR of 31.24% through 2034, investment flow analysis.

Electroiq / DemandSage / SQ Magazine / WifiTalents — Character.AI Statistics (2025–2026)

Monthly active users (20M+), 2-hr/day engagement, 373 min/week, 45% DAU/MAU rate, demographic breakdown (51.84% aged 18–24), 74% mobile traffic.

arXiv: Ebert et al. — Global Environmental AI Regulation (FAccT 2026)

Regulatory landscape analysis; proposed EU AI Act, Consumer Rights Directive, DSA amendments.

AI Metrix — Character.AI Environmental Impact Analysis (March 2026)

Per-message energy estimates (0.24–0.30 Wh/message using 2026 Blackwell B200 benchmarks); disclosure status audit as of March 2026.

EPA eGRID — US Grid Average Emissions Factor (2024)

US average grid emission factor used for CO₂ calculations: ~0.386 kg CO₂/kWh.

Methodology Notes

This report does not fabricate or extrapolate data without disclosure. Where verified figures were unavailable (e.g., Character.AI's own energy consumption), this report explicitly notes the absence of disclosure. Competitor session-duration comparisons for Claude and Gemini rely on traffic-proxy estimates clearly marked with † and "est." designations. TPU efficiency figures rely on Google's published specifications and independent technical analyses, not internal Character.AI data. Energy-per-chat calculations in Section 4 are presented as illustrative comparisons, not independently verified operational figures, and are clearly labelled as such. Market projections are sourced from named research firms and are attributed.

This report was produced by AllInsightsNews.net (June 2026) as a companion to "Is Character.AI Bad for the Environment?" (Oct 2025). It is intended for research and informational purposes. No content constitutes financial, legal, or investment advice. Energy calculations in Section 4

