THE SPARSE SUPERNOVA PODCAST

Episode: “The Watt Ceiling”

AI’s Energy Crisis — Who Runs Out of Power First?

FULL PRODUCTION SCRIPT — REVISED

Estimated runtime: 28–32 minutes | March 2026

Target audience: AI professionals · Climate experts · Business leaders

CHARACTER KEY

JOURNOBILL — Veteran tech journalist. Sharp, curious, plays devil’s advocate. Asks the questions the audience is thinking.

SPARSE BOB — CTO of Sparse Supernova. Edge-AI architect, systems thinker, carbon-conscious engineer. Brings data but keeps it human.

Based on the AI Energy Consumption Global Risk Assessment (March 2026), including Methodological Appendix.

COLD OPEN

[Ambient hum of a data centre — deep bass, fans whirring, rhythmic. Holds for 4 seconds. Fades slowly.]

SPARSE BOB: Here's a number that should keep every AI executive awake at night. Nine hundred and forty-five terawatt-hours. That's how much electricity data centres will need by 2030, according to the IEA's central scenario. The entire electrical output of Japan. All of it. And AI is the single biggest reason it's climbing.

[Sharp synth sting. Podcast theme drops in — upbeat, techy, confident. 8 seconds. Fades under.]

JOURNOBILL: You're listening to The Sparse Supernova Podcast, where we decode the signals from the noise. I'm JournoBill. Today's episode is called "The Watt Ceiling" and my guest builds AI systems at the edge, obsesses over every milliwatt, and has been loudly warning that the industry's energy bill doesn't add up. Bob, CTO of Sparse Supernova. Welcome back.

SPARSE BOB: Good to be here, Bill. Fair warning to your listeners: this one's going to be uncomfortable for anyone with a hundred-billion-dollar data centre on their balance sheet.

JOURNOBILL: Music to my ears. Let's get into it.

SEGMENT 1: HOW BIG IS THE PROBLEM?

Report Sections 1–2 — Global energy consumption, per-query costs, training energy. ~6 min.

JOURNOBILL: Paint me the picture, Bob. How much energy is AI burning right now, globally?

SPARSE BOB: 2024 numbers. Data centres worldwide consumed about 415 terawatt-hours of electricity. Around 1.5% of all global electricity. Sounds small.

JOURNOBILL: It does sound manageable.

SPARSE BOB: Until you see the growth curve. Twelve percent a year, five straight years. Four times faster than demand growth from every other sector on the planet combined. And AI servers specifically? Growing at 30% a year. Conventional servers? Nine percent. AI is eating the data centre from the inside out.

JOURNOBILL: AI's share of the total right now?

SPARSE BOB: The report's key metrics table puts it at 5 to 15% of data centre power today. Some aggregators — FlowingData, the IEA's own text — push the range to 11 to 20% depending on what you count. Either way, the direction is one way: by 2030, the IEA projects AI could be 35 to 50% of data centre electricity. AI-optimised data centre demand set to quadruple. And the stat that floored me: in their central AI-growth scenario, the United States will consume more electricity processing data by 2030 than manufacturing all its

energy-intensive goods combined. Aluminium. Steel. Cement. Chemicals. Data centres bigger than the lot.

JOURNOBILL: Data centres outstripping all US heavy industry? The IEA says that?

SPARSE BOB: The International Energy Agency's central scenario. Published, documented, cited.

[Brief reflective beat.]

JOURNOBILL: Bring this down to something I can hold. What does a single AI query cost in energy?

SPARSE BOB: The per-query data is murkier than anyone admits. Sam Altman has claimed a standard ChatGPT query uses about 0.34 watt-hours. An order of magnitude higher than a Google search. But OpenAI never published the methodology behind that number, so it's an indicative claim, not a verified constant. Since we wrote the first draft, an infrastructure-aware benchmarking paper estimates roughly four-tenths of a watt-hour for a short prompt in a GPT-4o class system — broadly the same order of magnitude — but still a range, not a constant.

JOURNOBILL: The bigger reasoning models?

SPARSE BOB: Third-party estimates — and I'm stressing estimates — suggest models like o3 might use 7 to 40 watt-hours per query. Up to a hundred times the basic model. Image generation? 20 to 40 times more than text. Video generation? A thousand to three thousand times more.

JOURNOBILL: So someone generating a five-second video clip with AI…

SPARSE BOB: Has potentially used energy in the same ballpark as a meaningful e-bike ride — from a single request. Andy Wu at Harvard Business School called generative AI "perhaps the most wasteful use of a computer ever devised." A calculator does 1+1 in one operation. A large language model does a trillion operations for the same answer. The compute overhead is extraordinary.

JOURNOBILL: Training the models themselves?

SPARSE BOB: GPT-4 is estimated at $50 to $100 million to train. The actual energy consumed? Never disclosed. Third-party reconstructions land north of 50 gigawatt-hours, but without a rock-solid primary citation. Next-generation frontier models being trained now could easily exceed 100 gigawatt-hours per run, based on current scaling estimates. Each run.

JOURNOBILL: And then DeepSeek came along.

SPARSE BOB: DeepSeek V3. The genuine counterpoint. Their technical report is admirably transparent: 2.788 million H800 GPU hours, about $5.6 million at cloud rates. Roughly 90 to 95% less compute cost than GPT-4. Important caveat: that's a cost and compute comparison, not a verified energy measurement. Actual kilowatt-hours depend on GPU power draw, utilisation, cooling efficiency, PUE — none disclosed. And the $5.6 million excludes prior research, ablation experiments, infrastructure costs. But the architectural efficiency — Mixture-of-Experts, FP8 training, latent attention — that's real. Frontier capability without city-scale electricity.

JOURNOBILL: Competitive AI without incinerating the grid?

SPARSE BOB: More efficient AI, yes. Whether the market rewards efficiency or just rewards whoever scales fastest… right now, mostly the latter.

[Transition beat. Brief.]

JOURNOBILL: Stick around — if the energy numbers are wild, wait until we get to the water.

SEGMENT 2: THE FORGOTTEN RESOURCE — WATER

Report Section 2 — Water consumption and material constraints. ~3 min.

JOURNOBILL: You've been banging on about water. Is this the story nobody's covering?

SPARSE BOB: Massively under-covered. Everyone talks watts. Almost nobody talks water. By 2030, global data centre water use is projected to hit 450 million gallons per day. The daily consumption of about 5 million people.

JOURNOBILL: Per day.

SPARSE BOB: Per day. A single large AI data centre can consume 300,000 to 5 million gallons a day for cooling alone. And two-thirds of US data centres sit in high water-stress regions. Virginia, Texas, Arizona — places already fighting over water for agriculture and drinking.

JOURNOBILL: Forward projection?

SPARSE BOB: By 2027, AI is projected to withdraw 1.1 to 1.7 trillion gallons of freshwater annually. Four to six times Denmark's entire annual water use. Texas alone may see an 870% increase in data centre water demand by 2030. Now — these are scenario projections, not exact forecasts. They're sensitive to siting decisions, cooling technology, and utilisation assumptions. But they give you the scale of what's coming.

JOURNOBILL: Eight hundred and seventy percent. One state.

SPARSE BOB: One state. Energy plus water, running in parallel, feeding each other, and largely invisible to the public and most investors.

[Transition. Globe-turning graphic sound begins to build.]

JOURNOBILL: That's the resource picture. Now let's look at where the grid actually breaks.

SEGMENT 3: THE MAP — WHO RUNS OUT OF POWER FIRST?

Report Section 3 — Territorial shortfalls: Virginia, Texas, Ireland, Europe, Asia-Pacific, infrastructure. ~7 min.

[Globe-turning graphic sound lands. Brief whoosh.]

JOURNOBILL: Around the world. Where are the real breaking points?

SPARSE BOB: Northern Virginia. Data Center Alley. Biggest concentration of data centres on the planet. Already consuming roughly 26% of Virginia's electricity.

JOURNOBILL: A quarter of the state's electricity. Data centres alone.

SPARSE BOB: Data centres alone. PJM, the regional grid operator, has seen capacity market prices go from about $29 per megawatt-day to $329. More than tenfold. Directly attributable to data centre demand. The projected grid upgrade bill? $90 billion. Northern Virginia is approaching saturation — new projects getting pushed south and west because the grid is full.

JOURNOBILL: Bloomberg numbers?

SPARSE BOB: BloombergNEF forecasts data centre capacity in the PJM grid at 31 gigawatts by 2030. New generation capacity expected to come online in the same period? 28.7 gigawatts. Demand growing faster than supply. Collision course.

JOURNOBILL: Texas?

SPARSE BOB: ERCOT — Texas runs its own independent grid. Reserve margins could go dangerously thin after 2028. They've already reacted: Senate Bill 6 means any load over 75 megawatts faces a new interconnection process. Disclosure, financial commitments, proof you can cover transmission infrastructure costs. No more plugging in a gigawatt campus and hoping for the best. And the water — 870% projected increase by 2030.

[Brief contemplative beat.]

JOURNOBILL: Across the Atlantic. Ireland.

SPARSE BOB: Ireland is the cautionary tale the whole industry should study. CSO Ireland data: data centres consumed 22% of total national metered electricity in 2024. Up from 5%

in 2015. Less than a decade from rounding error to the single largest electricity consumer in the country, overtaking all urban households combined.

JOURNOBILL: Twenty-two percent nationally. Dublin specifically?

SPARSE BOB: Extreme. Oeko-Institut analysis, cited by Carbon Brief, estimates data centres consume approximately 79% of Dublin's electricity. I want to be rigorous: that figure hasn't been independently replicated and the full methodology isn't public. In our report's methodological appendix, we classify it Tier 3 — a single-study estimate, useful for scale, not for precise calibration. But it's plausible. EirGrid confirms almost all contracted data centre capacity sits in the Greater Dublin region.

JOURNOBILL: Seventy-nine percent. A European capital's electricity. That's not a data centre cluster — that's a power vampire draining a city.

SPARSE BOB: Ha. Vivid, Bill. And this particular vampire triggered amber power alerts, a moratorium on new data centres until 2028, and a complete rethink of grid connection policy. The CRU only lifted the moratorium conditionally in late 2025. New rules: bring your own dispatchable generation or battery storage matching full import capacity, and source 80% from renewables. You want a data centre in Ireland? Bring your own power station.

JOURNOBILL: EirGrid projects 31 to 32% of national electricity by 2027?

SPARSE BOB: They do. And Ireland has legally binding carbon budgets under the Climate Action Plan. There's a real scenario where data centres make it impossible for Ireland to meet its Paris obligations.

JOURNOBILL: Rest of Europe?

SPARSE BOB: European data centre demand projected to grow about 70% by 2030 — more than 45 additional terawatt-hours. The EU responded with the Energy Efficiency Directive requiring reporting above 500 kilowatts, and the EU AI Act mandating energy disclosure for general-purpose AI models. Enforcement mechanisms remain unclear.

JOURNOBILL: Asia-Pacific?

SPARSE BOB: China and the US together account for nearly 80% of global AI electricity consumption growth to 2030. China's mandating PUE of 1.25 for large data centres and pushing state-owned facilities to 100% renewables by 2030. Southeast Asia is the emerging hotspot — Singapore, southern Malaysia — demand expected to more than double. Japan's striking: data centres could account for more than half the country's entire electricity demand growth.

JOURNOBILL: Half of Japan's demand growth.

SPARSE BOB: Half. And Malaysia could see a fifth of its total electricity growth consumed by data centres. Real trade-offs: power AI, or power industrialisation and households.

[Pause beat.]

JOURNOBILL: And the infrastructure to deliver all this?

SPARSE BOB: Can't keep up. A data centre goes operational in two to three years. The grid behind it takes 5 to 10 years to plan, permit, and build. Fundamental temporal mismatch. High-voltage transformers that used to have six-month lead times are now at three to four years. The IEA estimates roughly 20% of planned data centre projects face delay risk purely because the grid can't match the pace. Grid access has become the primary bottleneck globally. Ahead of land, ahead of capital.

JOURNOBILL: The binding constraint.

SPARSE BOB: The binding constraint.

[Transition. Cash register sound begins ominously.]

JOURNOBILL: If you think that's uncomfortable, wait until you see the money.

SEGMENT 4: THE MONEY — DOES THE INVESTMENT THESIS SURVIVE?

Report Section 4 — Scale of investment, circular financing, enterprise returns, energy-returns contradiction, chip obsolescence. ~7 min.

[Cash register sound effect lands. Slightly ominous.]

JOURNOBILL: How much is going into AI infrastructure right now?

SPARSE BOB: The four biggest hyperscalers — Microsoft, Google, Amazon, Meta — on track to spend over $325 billion collectively this year. A hundred billion more than expected at the start of 2025. Add neoclouds, sovereign clouds, private clouds, and you're at $600 billion for 2025 alone. Could hit a trillion by 2027 or 2028.

JOURNOBILL: A trillion dollars a year on data centres.

SPARSE BOB: And the macro context: AI capex accounted for over one percentage point of US GDP growth in the first half of 2025. It overtook the consumer as the primary driver of American economic growth. The share of GDP going to AI investment is nearly a third larger than internet investment during the dot-com bubble. The Manhattan Project was 0.4% of GDP. Apollo was similar. This dwarfs both.

JOURNOBILL: "Unprecedented" doesn't cover it.

SPARSE BOB: Not close.

JOURNOBILL: The financing structures. Walk me through what worries you.

SPARSE BOB: OpenAI committed $300 billion with Oracle for computing power. $60 billion a year, against projected revenues of $13 billion. CNBC reporting suggests Oracle expects to lose considerable sums on the deal. Meanwhile, OpenAI took a 10% stake in AMD. Nvidia invested $100 billion in OpenAI. Microsoft is a major shareholder in OpenAI but also a major customer of CoreWeave, where Nvidia holds significant equity. And Microsoft accounts for almost 20% of Nvidia's revenue.

JOURNOBILL: Everybody's investor, customer, and supplier simultaneously.

SPARSE BOB: Exactly the pattern. Howard Marks at Oaktree called it "bubble-like behaviour." Yale published a piece called "This Is How the AI Bubble Bursts" mapping the circular structure. When everyone depends on everyone else, a disruption to any single node cascades through the whole web.

JOURNOBILL: What could trigger the cascade?

SPARSE BOB: Energy. If the grid can't deliver, data centres can't operate. If data centres can't operate, revenue projections don't materialise. If revenue doesn't materialise, debt can't be serviced. The entire structure assumes abundant, available power. Pull that assumption out and the arithmetic collapses.

JOURNOBILL: The MIT report.

SPARSE BOB: NANDA. MIT Media Lab, August 2025. "Despite $30 to $40 billion in enterprise investment into generative AI, 95% of organisations are getting zero return." Important nuance: that metric measures measurable business transformation — deployed pilots generating revenue or operational savings. Doesn't mean the technology failed in every deployment. Some organisations are still in pilot stages. But the gap between capital going in and value coming out is enormous. And it matters here because the demand assumptions behind all this infrastructure rest on explosive, sustained AI adoption. If 95% of enterprise buyers aren't seeing returns, the demand curve might be far softer than supply is building for.

JOURNOBILL: Sam Altman said he thinks there's a bubble.

SPARSE BOB: He did. So did Ray Dalio at Bridgewater. The Bank of England warned infrastructure requirements might be too high to be met and that investors weren't properly cautioned about crash risk. When the leading AI CEO, the world's most prominent hedge fund manager, and a central bank all flag the same concern…

JOURNOBILL: Devil's advocate. JPMorgan says no bubble. BlackRock says valuations are reasonable. Jerome Powell distinguished this from dot-com.

SPARSE BOB: And there's genuine substance there. S&P 500 IT Index at 30 times forward earnings — elevated, well below the 55 times at dot-com peak. Today's tech giants have real revenues, real cash flows. Microsoft, Meta, Amazon, Google can absorb losses.

The risk concentrates in the layer below: OpenAI, xAI, CoreWeave, Oracle's data centre bets. Debt-financed commitments without hyperscaler cash flows behind them. That's where energy constraints bite first.

JOURNOBILL: Chip obsolescence.

SPARSE BOB: The silent killer. AI chips become technically obsolete in 10 to 12 months. Physical stress under high utilisation means replacement every two and a half years or so. Think about what that means for infrastructure planning. Imagine laying railroad tracks in 2026, only to rip them up in 2028 because "faster trains" need a completely different gauge. Then again in 2030. And again. That's AI hardware. Every generation demands new silicon, new cooling, new power. Capital deployed today may need complete replacement before it generates returns. A perpetual treadmill — and every cycle of the treadmill consumes more energy than the last.

SEGMENT 5: CARBON, CLIMATE, AND CALLING OUT GREENWASH

Report Section 6 — Carbon/Paris alignment, Jevons Paradox, disclosure framework. ~4 min.

JOURNOBILL: Climate. Your report draws a hard line between energy consumption and carbon emissions. Why does that distinction matter?

SPARSE BOB: Because it gets muddled constantly, and the muddling serves people who don't want scrutiny. Terawatt-hours and kilograms of CO2 are different things. If every data centre ran on genuinely clean energy, the resource constraint would remain but the climate problem would shrink dramatically. Most data centres don't run on clean energy.

JOURNOBILL: IEA on actual emissions?

SPARSE BOB: Central scenario: data centre emissions reach about 1% of global CO2 by 2030. Growing from 220 megatonnes in 2024 to 300 to 320 megatonnes by 2035. High scenario pushes to 500 megatonnes. One percent sounds contained, until you realise data centres are one of only three sectors where emissions are set to grow — alongside road transport and aviation. Everything else is supposed to be heading the other direction.

JOURNOBILL: Harvard finding on carbon intensity?

SPARSE BOB: Harvard's T.H. Chan School of Public Health found data centre electricity is 48% more carbon-intensive than the US average. Because they draw 24/7 baseload power. They can't simply switch off when the sun goes down or the wind drops. Even with renewable PPAs on paper, the actual electrons flowing through the meter are dirtier than average.

JOURNOBILL: Greenwash.

SPARSE BOB: The press releases write themselves. "100% renewable match!" In practice? Renewable energy certificates covering annual consumption. Not clean electrons

flowing through the facility 24 hours a day, 365 days a year. The grid doesn't work that way. Reporting only market-based emissions while hiding the location-based reality is greenwash. Full stop.

JOURNOBILL: What should proper reporting look like?

SPARSE BOB: Four metrics. One: kilowatt-hours per thousand inferences — standardised energy intensity, with architecture and hardware specs attached. Two: both location-based and market-based CO2 emissions, side by side. The location-based figure shows what actually came out of the grid. Three: a carbon budget benchmarked to a 1.5-degree pathway — quantified, year-by-year trajectory. Four: absolute emissions targets alongside intensity targets.

JOURNOBILL: Why do absolute targets matter alongside intensity?

SPARSE BOB: The whole game. Make AI 50% more efficient per query, then run ten times more queries — total emissions up fivefold. Intensity improvements that get reinvested into volume expansion are not climate progress. That's Jevons Paradox in action.

JOURNOBILL: Jevons. For anyone who doesn't know the name.

SPARSE BOB: William Stanley Jevons. 19th-century economist. Observed that when coal engines became more efficient, total coal consumption rose, not fell. Cheaper energy made more uses economical. Precisely what's happening with AI compute now. Every efficiency gain gets absorbed by scaling. The IEA's own data: 12% annual data centre demand growth, sustained, despite continuous efficiency improvements. If efficiency gains aren't captured as absolute emission reductions — if they're simply reinvested into growth — those additional emissions need separate accounting and separate offsetting. That's governance territory.

JOURNOBILL: Engineering alone won't solve this.

SPARSE BOB: Governance solves this. Engineering enables it.

SEGMENT 6: THE WAY FORWARD — WHO WINS?

Report Section 7 — Conclusions, sovereign/edge architectures, model efficiency, nuclear, governance. ~4 min.

JOURNOBILL: OK Bob, we've thoroughly depressed everyone. Who actually survives The Watt Ceiling?

SPARSE BOB: The edge wins. I know that sounds self-serving from a company called Sparse Supernova, but the logic holds regardless of who builds it. Systems designed for energy efficiency, governance-by-design, running inference at the network edge instead of in a centralised hyperscale barn. Those architectures are built for constrained energy. When

the grid can't deliver another gigawatt to Virginia, companies running efficient models on distributed infrastructure keep operating.

JOURNOBILL: What else?

SPARSE BOB: Model architecture choices. University of Rhode Island research shows energy consumption varies by 70 times or more between model options for equivalent tasks. Choosing the right model for the job — not defaulting to the biggest — is one of the highest-leverage decisions a technology leader can make. DeepSeek proved frontier performance with radically less compute. The industry should be racing toward efficiency, not toward raw scale.

JOURNOBILL: Nuclear? Small modular reactors?

SPARSE BOB: The IEA expects the first SMRs online around 2030. Nuclear has a genuine role: dispatchable, low-carbon, runs 24/7, matches data centre load profiles perfectly. Three Mile Island getting revived for data centre power tells you how desperate the near-term situation is. But SMRs are a 2030s solution. The constraint is 2026.

JOURNOBILL: The 2026 answer?

SPARSE BOB: Discipline. Don't build capacity you can't power. Locate where grid capacity exists today. Operate flexibly — the IEA notes AI data centres are ten times more capital-intensive than aluminium smelters, so curtailment hurts, but so does sitting idle when the grid can't deliver. On the governance side: mandatory energy and carbon disclosure, Paris-aligned reporting, and regulators willing to say no. Ireland said no. More will follow.

JOURNOBILL: Push back: if you're an investor, building fastest wins the market. Hesitation is the risk.

SPARSE BOB: And that's exactly the logic that inflates bubbles. I build AI systems for a living. The technology matters. But building on the assumption that energy is infinite and available everywhere is magical thinking. Not every AI company can win the energy race — the physics won't allow it. Some will commit hundreds of billions to facilities that can't get grid connections, face regulatory curtailment, or get stranded when the next GPU generation makes their hardware obsolete. The winners will be the ones who built energy awareness into their strategy from day one. Not as an afterthought. From day one.

JOURNOBILL: Who specifically are you worried about?

SPARSE BOB: Without naming names: the debt-financed layer. Companies without hyperscaler cash flows carrying hyperscaler-scale commitments. Companies whose revenue depends on circular financing — where everyone's income depends on everyone else's spending. If demand grows even slightly slower than projected — and 95% of enterprises aren't seeing returns yet — the arithmetic gets difficult fast.

JOURNOBILL: So, Bob. Listeners walk away from this episode. One action item.

SPARSE BOB: Audit your AI stack's energy footprint. Today. Know what you're consuming, know where the electrons come from, know what it costs in carbon. Because tomorrow, the grid might not be there to ask nicely.

CLOSE

[Theme music begins to fade in softly underneath.]

JOURNOBILL: Bob, the one-liner. A CEO, a climate minister, and an AI engineer are all listening.

SPARSE BOB: Energy is not a footnote to the AI strategy. Energy is the AI strategy. If you don't have a credible plan for how your systems get powered — sustainably, at scale, in the specific territories you operate in — you don't have a strategy. You have a hope.

JOURNOBILL: Sparse Bob from Sparse Supernova. Thanks for a characteristically uplifting conversation.

SPARSE BOB: Ha! Any time, Bill. Tell your listeners to check their electricity bills.

[Theme music up. Plays for 6 seconds.]

JOURNOBILL: That's it for this episode of The Sparse Supernova Podcast. The full AI Energy Consumption Global Risk Assessment — including the new Methodological Appendix explaining how we classify every stat — is linked in the show notes. Every source referenced today: IEA, IMF, Goldman Sachs, MIT NANDA, Carbon Brief, all of it. If this episode made you think, send it to someone who needs to hear it. I'm JournoBill. See you next week.

[Theme music swells. Fades to silence over 5 seconds. END.]

PRODUCTION NOTES

REVISION NOTES (v2)

Incorporates review feedback from independent fact-check passes. Changes from v1:

- *Per-query comparison: "roughly ten times" → "an order of magnitude higher" (more defensible)*
- *Frontier training >100 GWh: softened to "could easily exceed" (projection, not disclosed fact)*
- *IEA heavy-industry comparison: explicit "central AI-growth scenario" attribution added*
- *AI share of DC power: harmonised 5–15% (table) with 11–20% (aggregator range), both cited*

- *Water projections: explicit scenario-modelling caveat added*
- *Dublin 79%: tied to Tier 3 classification in Methodological Appendix*
- *Chip obsolescence: expanded railroad analogy for listener engagement*
- *Segment transitions: listener hooks added for pacing*
- *Close: actionable takeaway added ("Audit your AI stack's energy footprint")*
- *Show notes: Methodological Appendix from Report v3 now referenced*
- *Closing one-liner extended: "you don't have a strategy, you have a hope"*

RUNTIME TARGET: 28–32 minutes

Cold Open: 1 min | Seg 1 — Scale: 6 min | Seg 2 — Water: 3 min

Seg 3 — Map: 7 min | Seg 4 — Money: 7 min | Seg 5 — Carbon: 4 min

Seg 6 — Way Forward: 4 min | Close: 1 min

VOICE DIRECTION

JOURNOBILL: Warm, confident, slightly wry. Surprise reactions earned, not performed. Seasoned broadcaster who enjoys the subject. Hooks at segment transitions should feel natural, like a presenter who genuinely wants to keep listeners.

SPARSE BOB: Authoritative, accessible. Numbers grounded in human terms. Dry humour lands when he drops it. Gets passionate on governance and climate. Audience should trust him implicitly.

SFX: Minimal. Data centre hum at open sets the scene. Theme used sparingly. Cash register in Seg 4 adds levity before heavy content. Globe-turning sound in Seg 3 for geographic transitions. No applause, laugh tracks, forced transitions.

FACT-CHECK ANCHORS

All statistics traceable to Report (March 2026) footnotes [1]–[26] and Methodological Appendix (Tier 1/2/3 classification).

Per-query figures: Tier 3 (estimates, explicitly flagged). Dublin 79%: Tier 3 (Oeko-Institut/Carbon Brief, single-study). DeepSeek: compute cost, not verified energy (explicitly flagged). MIT NANDA 95%: measurable business transformation metric (nuance preserved). IEA projections: Tier 2 (central scenario, documented methodology).

SHOW NOTES LINKS

IEA Energy and AI (2025) | IMF WP/2025/081 | Goldman Sachs AI/DC Power Demand | BloombergNEF AI and the Power Grid | DeepSeek V3 Technical Report (arXiv:2412.19437v1) | Yale Insights "This Is How the AI Bubble Bursts" | MIT NANDA "The GenAI Divide" | CSO Ireland data centre electricity releases | Carbon Brief "AI data centre energy in context" | Harvard T.H. Chan School carbon intensity study | Jegham et al. "How Hungry is AI?" (arXiv:2505.09598v6) | Full AI Energy Consumption Global Risk Assessment (March 2026, incl. Methodological Appendix)