AutoPresent

The Kilowatt Crunch: Balancing AI Ambitions with Net-Zero Realities

Strategic Alignment of Compute Infrastructure and ESG Goals

January, 2026

Discussion Topics

Integrated governance aligns AI growth with net-zero, turning the kilowatt crunch into competitive advantage.

Topic	Key Points
The Strategic Dilemma	• AI compute demand projected to increase 160% by 2030 while net-zero targets require 50% emissions reduction • This creates an untenable choice: abandon AI leadership or ignore ESG commitments • The 'Kilowatt Crunch' is a board-level risk, not an operational issue
Why Superficial Strategies Fail	• Offloading to cloud merely hides emissions (hyperscaler emissions up 30-48%) • Offsetting becomes financially unsustainable within 5 years • Optimizing in isolation is too slow (tech breakthroughs 5-10 years out vs. 24-month regulatory pressure)
The Three-Pillar Solution	• Optimize Compute: Hardware refresh (high CAPEX, 24+ month lead times) and software efficiency (faster but talent-constrained) • Decarbonize Supply: PPAs and renewable sourcing limited by grid capacity in key hubs • Integrate Governance: The critical enabler that aligns technical and energy levers
Governance as Game Changer	• Establish a 'Compute & Carbon Board' to align CIO and CSO incentives • Implement carbon-per-workload reporting and carbon-aware budgeting • Overcome cultural barriers between tech and sustainability teams
Path to Competitive Advantage	• Leverage scale to demand carbon transparency from cloud providers • Attract top talent (67% prefer sustainable employers) • Achieve 10% lower cost of capital through ESG leadership

Source: International Energy Agency (IAE / IEA) estimates on AI, data centres and electricity demand, Hyperscaler emissions data (major cloud providers such as AWS, Microsoft Azure, Google Cloud), PPAs (Power Purchase Agreements) and renewable energy sourcing disclosures from utilities and corporate buyers, AutoPresent Analysis

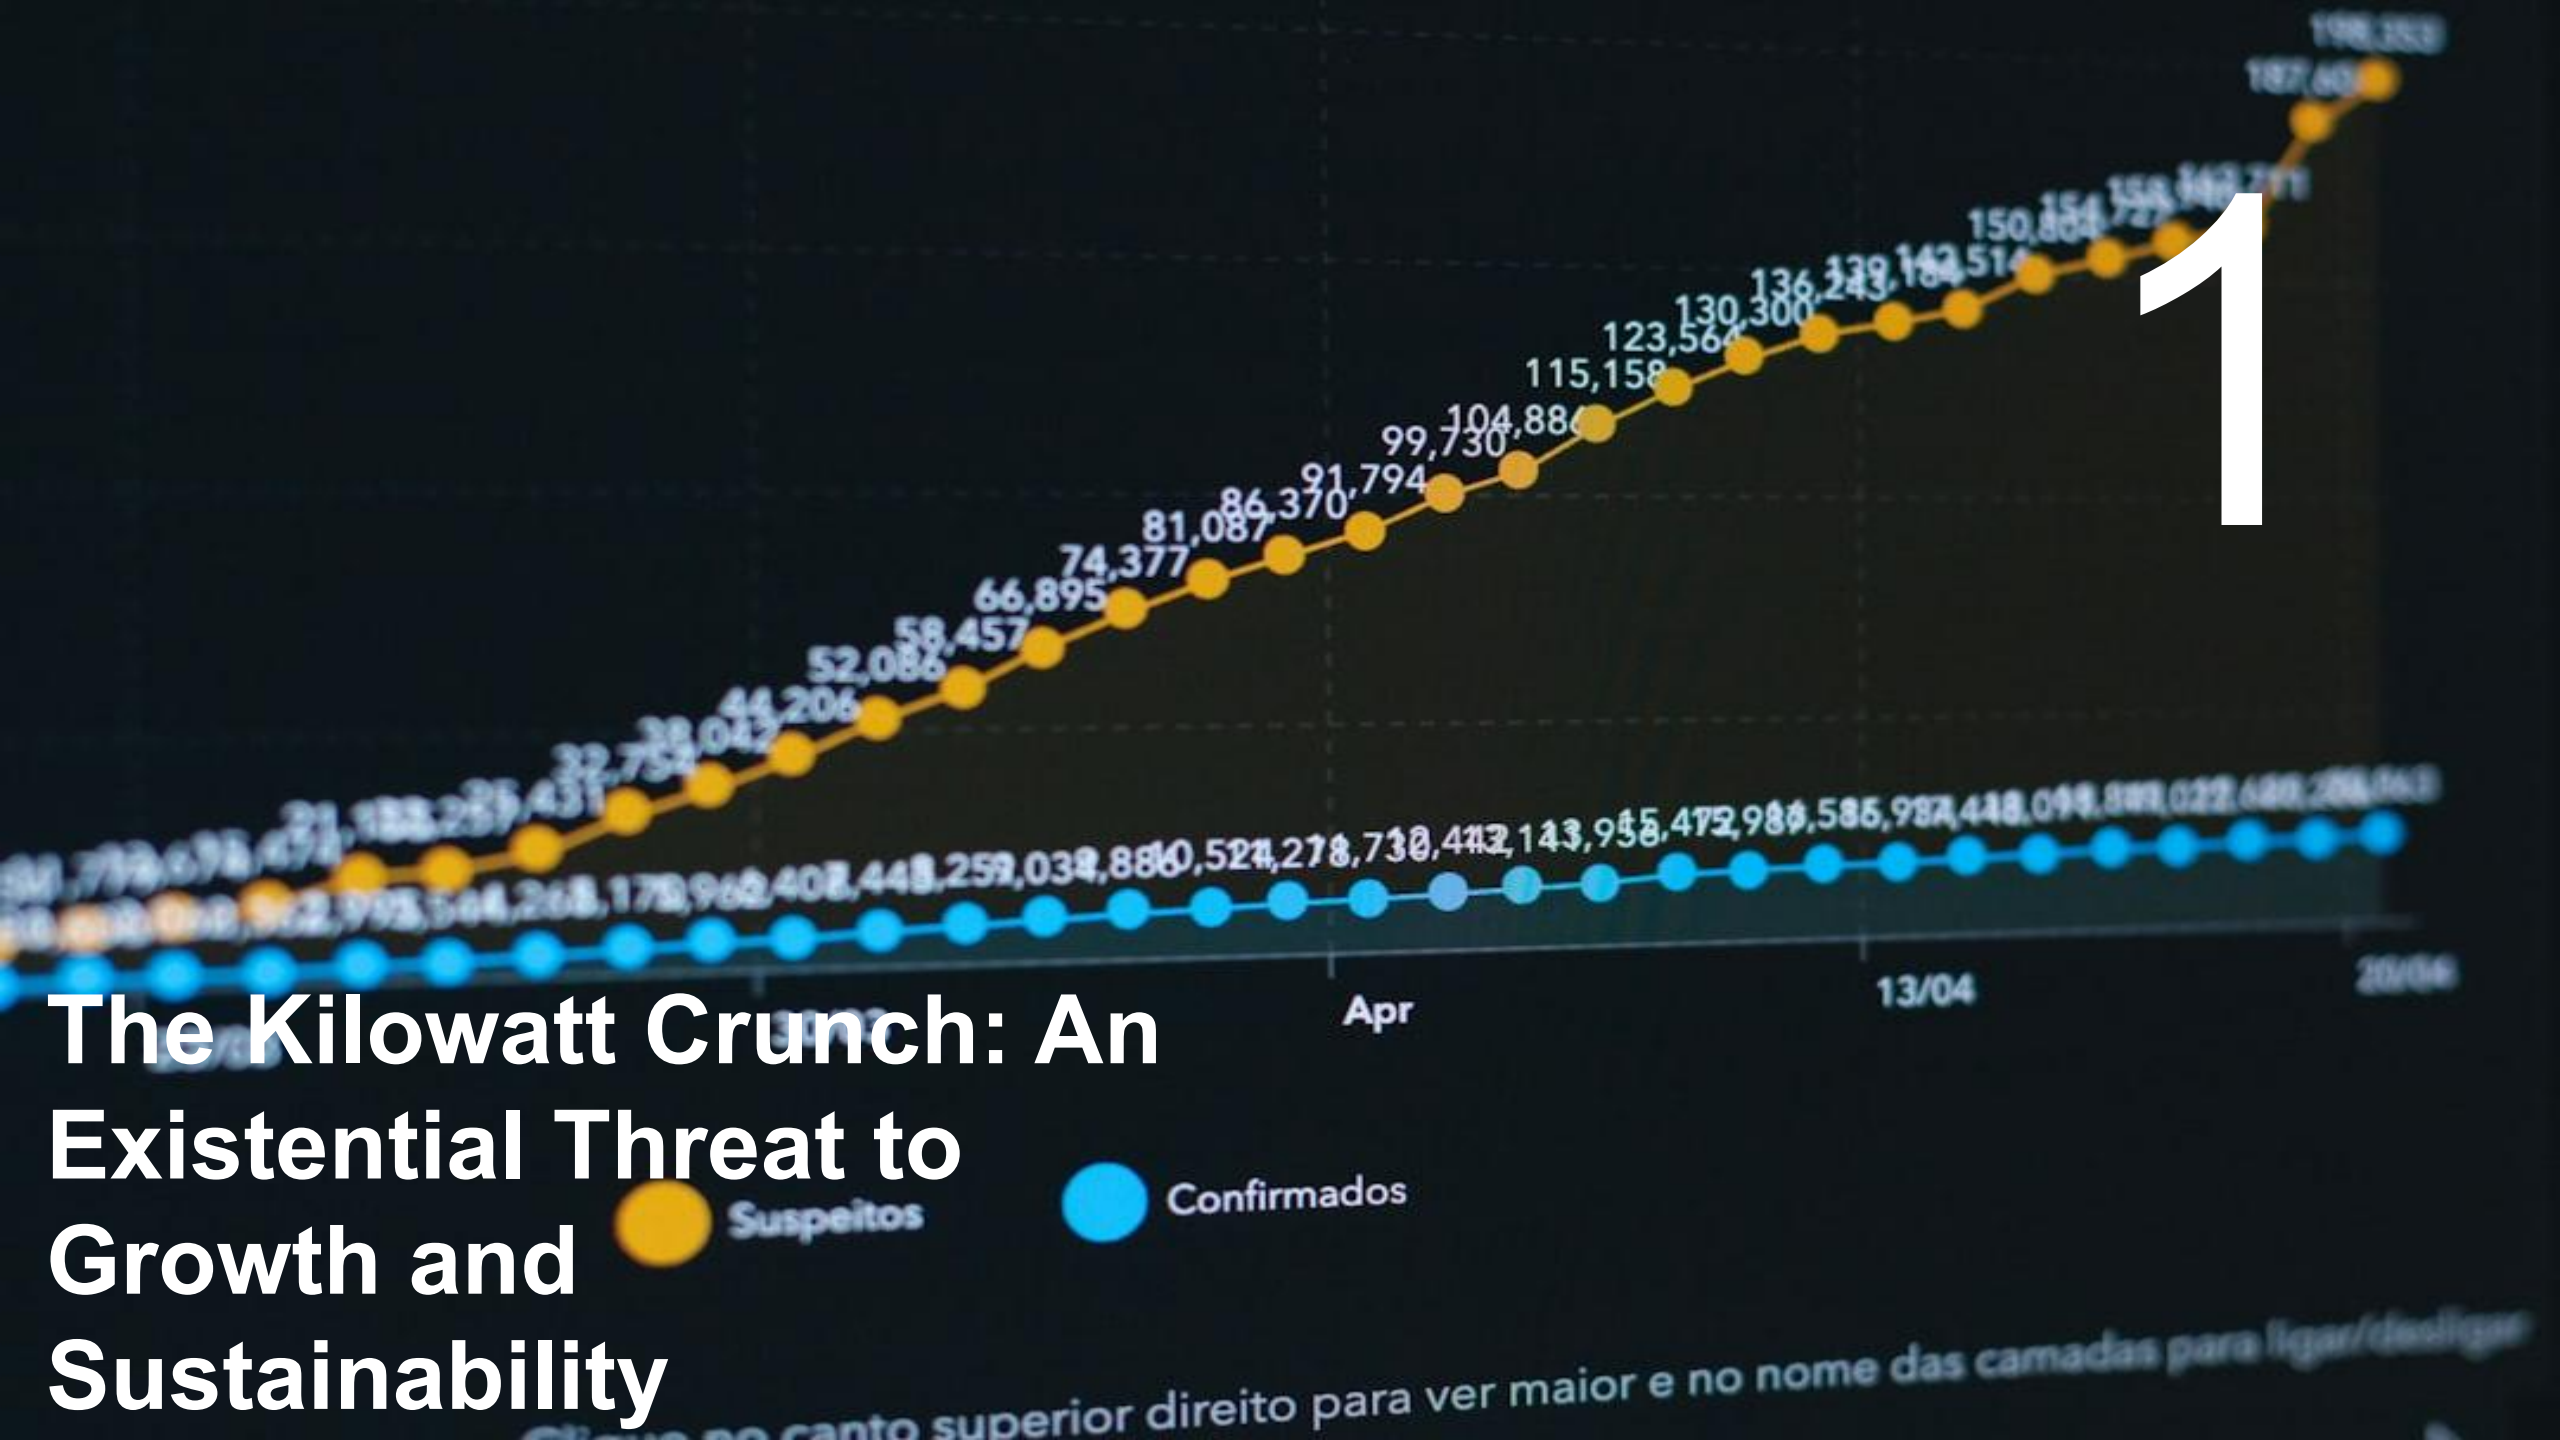

The Kilowatt Crunch: An Existential Threat to Growth and Sustainability

AI's exponential compute demand is on a direct collision course with flatlining corporate carbon budgets, creating an existential strategic risk for any data-intensive organization.

Structural Collision: +165% AI Demand Surge vs. Mandatory Absolute Scope 2 Reductions

Data Center Electricity Demand (TWh) | Year (2023-2030)

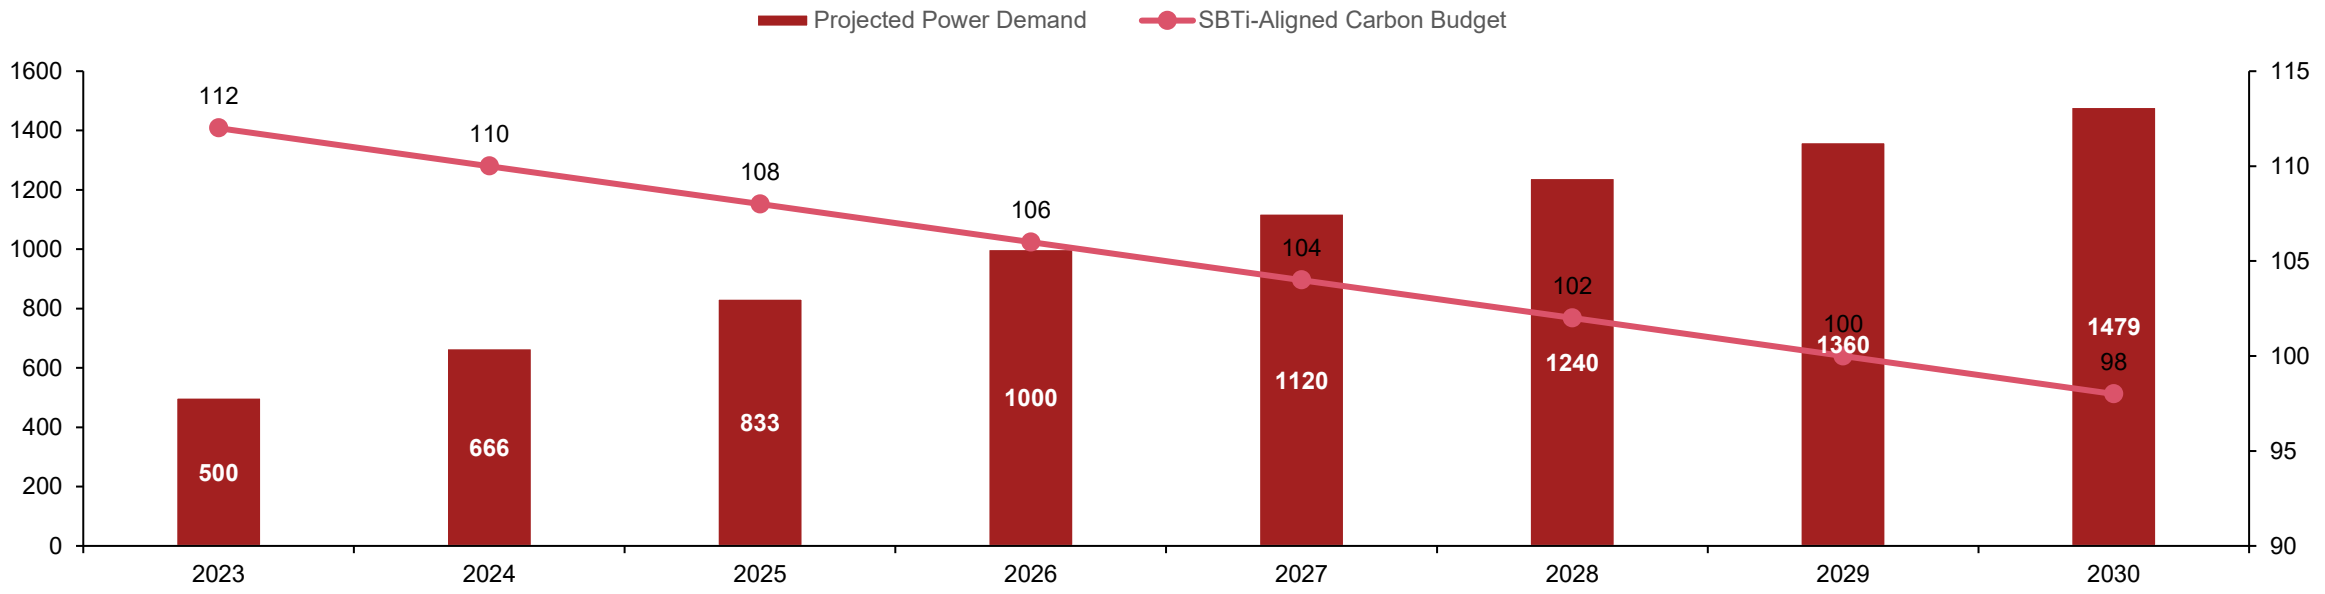

ESG Covenants & Capital Access:
Unchecked AI scaling triggers a breach of SBTi v2 mandates by 2026, risking institutional divestment and regulatory intervention. This creates a binary trade-off: sacrifice AI-driven growth or forfeit ESG-linked capital access and license to operate.

Source: IEA Study on Data Center Energy Consumption, AI Energy Consumption: Statistics from Key Sources [2026], United States Data Center Energy Usage Report, 2024, US data centers' energy use amid the artificial intelligence boom (Pew Research Center), AI has high data center energy costs — but there are solutions (MIT Sloan), AI: Five charts that put data-centre energy use and emissions into context (Carbon Brief), Can US infrastructure keep up with the AI economy? (Deloitte), 2026 Predictions: AI Sparks Data Center Power Revolution (Data Center Knowledge), Power for AI: Easier Said Than Built (BloombergNEF), Powering the US Data Center Boom: Why Forecasting Can Be So Difficult (World Resources Institute), Electric Power Research Institute (EPRI) Data on Regional Data Center Power Consumption, Goldman Sachs Research on Data Center Power Demand Forecasts, May 2025 Patterns Study on Generative AI Model Energy and Emissions (369-model cohort), Greenly February 2025 Estimate on AI Carbon Footprint, SBTi Corporate Net-Zero Standard v2 (March 2025), IEA Energy Attribute Certificates (EACs) and Renewable Energy Certificate (REC) Guidance, AutoPresent Analysis

To quantify this collision, AI workloads are projected to increase data center energy consumption by over 160% by 2030, while most corporate net-zero targets require a 50% reduction in scope 1 & 2 emissions.

AI Energy Growth vs. Net-Zero Targets: The Widening Sustainability Gap

Index (2024 = 100), Timeline

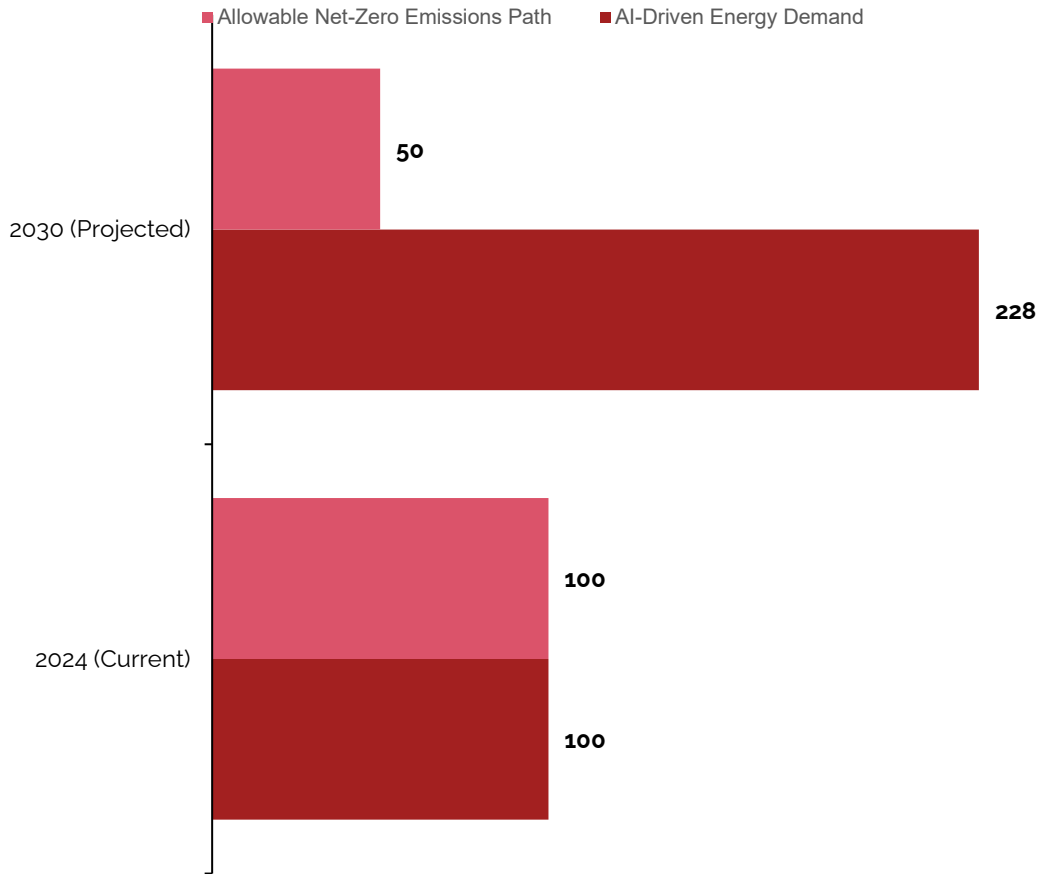

Structural Divergence: Projected AI Load Growth vs. Mandated Emissions Reductions

a — AI intensity scales non-linearly vs. traditional compute

- GenAI prompts use 10-100x more energy than standard search; AI to reach ~9% of total U.S. grid demand by 2030.
- Inference workloads operating at suboptimal utilization further inflate intensity per compute unit.

b — Net-zero standards are tightening to exclude offsets

- SBTi 2025 mandates >90% absolute reduction in Scope 1/2; reliance on unbundled RECs/offsets is increasingly restricted.
- Decarbonization now requires physical grid alignment, which lags behind the 18% CAGR of AI power demand.

c — Capital alone cannot decouple growth from emissions

- Microsoft's $6B renewable investment has not reversed its rising AI-driven energy intensity (2025 data).
- The 'Sustainability Gap' forces a strategic trade-off between AI deployment velocity and ESG integrity.

Source: IEA, Goldman Sachs Research, SBTi Corporate Net-Zero Standard, Microsoft, MIT News, Deloitte Technology, Media and Telecommunications (TMT) Predictions, Goldman Sachs report "AI to drive 165% increase in data center power demand by 2030", World Economic Forum article "How data centres can avoid doubling their energy use by 2030", Pew Research Center article "What we know about energy use at US data centers amid the AI boom", Electric Power Research Institute paper on U.S. data center electricity demand (2024), Boston Consulting Group analysis on U.S. data center electricity demand projections, Lawrence Berkeley National Laboratory study on data center demand, McKinsey report on data center power demand (referenced in World Economic Forum article), AutoPresent Analysis

This widening "Sustainability Gap" will force an impossible choice: abandon AI leadership to meet ESG goals, or ignore ESG commitments to pursue AI capabilities.

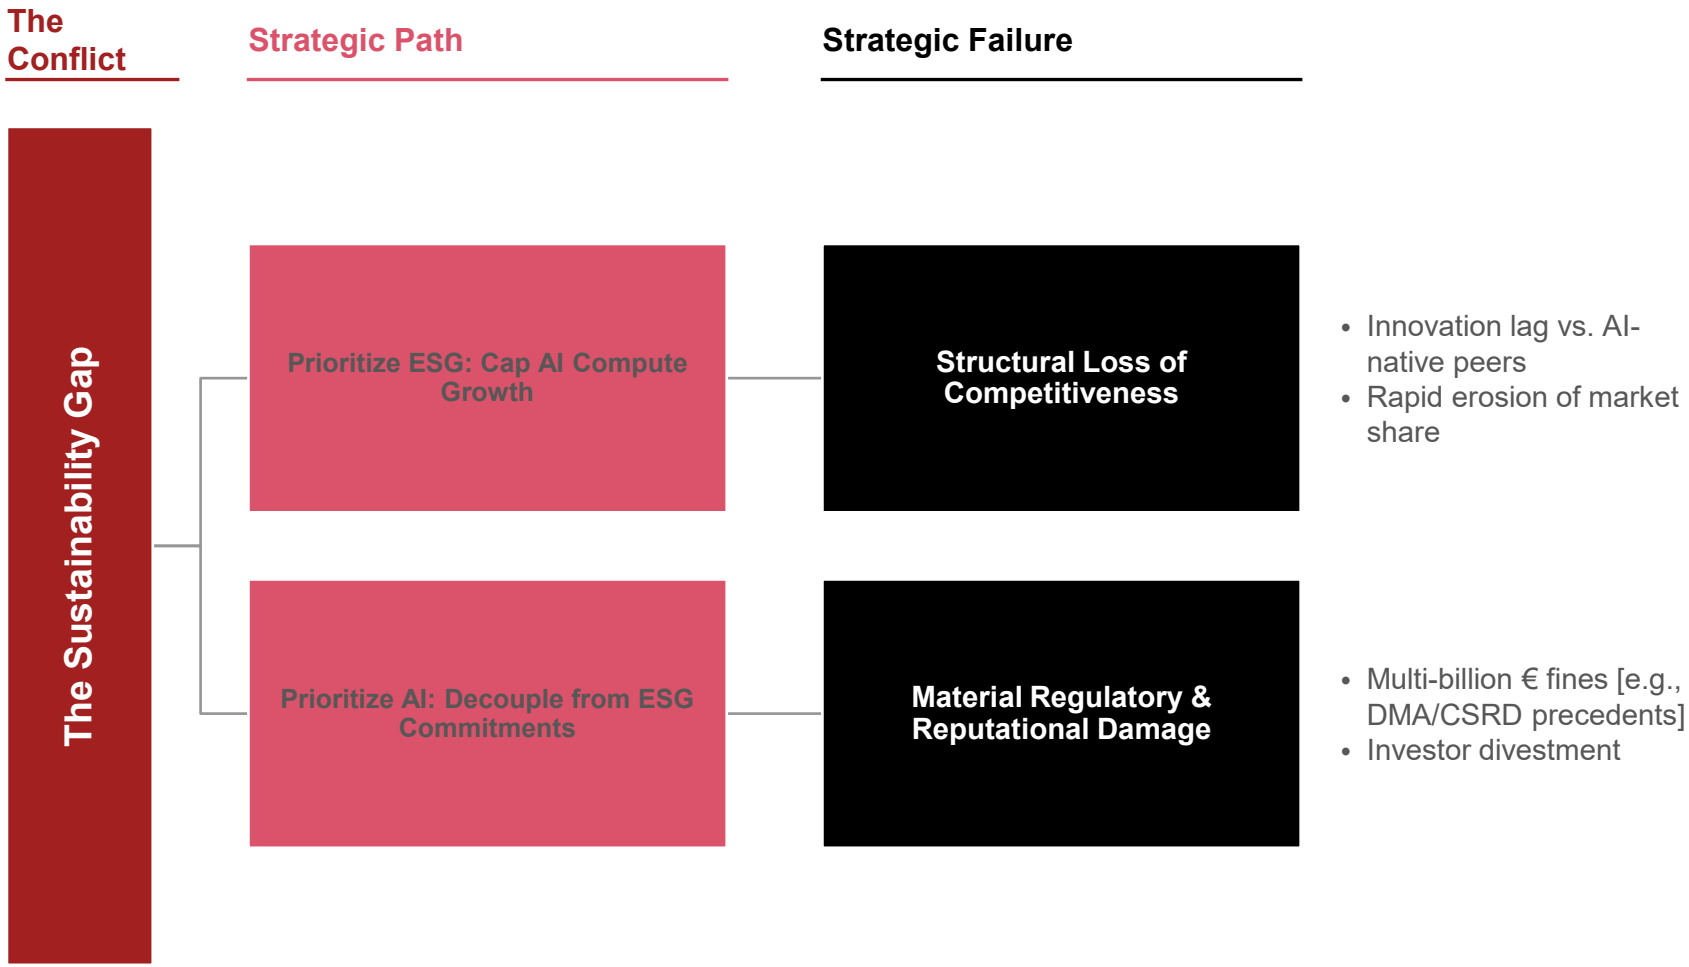

The Mathematical Impossibility of Status-Quo Efficiency

- The Growth Surge: Global DC power demand to rise 165% by 2030, driven by an AI workload shift from 15% to 50% of total consumption
- The Efficiency Lag: 1–2% annual legacy efficiency gains are mathematically unable to offset a 1,050 TWh consumption spike by 2026
- The Grid Constraint: DC demand potentially reaching 10% of total U.S. generation by 2035, forcing regulatory caps on unmanaged growth
- The Zero-Sum Result: Incremental PUE tuning cannot bridge this gap; AI ambitions now exceed the physical limits of current Net-Zero roadmaps

Source: International Energy Agency (IEA) report on data center electricity consumption, Goldman Sachs Research forecast on data center power demand, 451 Research analysis on grid-power demand from data centers, BloombergNEF regional data center power demand modeling, Pew Research Center article "What we know about energy use at US data centers amid the AI boom", Electric Power Research Institute paper on U.S. data center electricity demand, World Resources Institute article "Powering the US Data Center Boom: Why Forecasting Can Be So ...", Lawrence Berkeley National Laboratory study on data center demand, Boston Consulting Group projection of U.S. data center electricity demand to 2030, Data Center Frontier article summarizing IEA study on AI and cryptocurrency doubling data center energy consumption by 2026, Lightwave (fiber-optic telecommunications and networking publication), AI Multiple article "AI Energy Consumption: Statistics from Key Sources [2026]", U.S. Congressional Research Service report "Data Centers and Their Energy Consumption", European Commission article "In focus: Data centres – an energy-hungry challenge", DMA (Digital Markets Act) enforcement cases – Apple €500M fine, DMA (Digital Markets Act) enforcement cases – Meta €200M fine, CSRD (Corporate Sustainability Reporting Directive), AutoPresent Analysis


2
Why Superficial Strategies Are Doomed to Fail

Because this choice is untenable, most organizations are defaulting to one of three superficial strategies: Offload to the Cloud, Offset with Credits, or Optimize in Isolation.

Offload: The 'Scope 3' Transparency Trap

- **Accounting Arbitrage:**
  - Shifting to hyperscalers masks absolute growth
  - AI is projected to drive 35-50% of data center power by 2030 as provider emissions continue to rise despite clean energy buys.

Offset: The High-Integrity Credibility Gap

- **Integrity Risk:**
  - Large-scale offsets fail emerging 'high-integrity' criteria (SBTi/ICVCM)
  - 2025 EU EED mandates shift the focus from relative offsets to absolute energy and environmental accountability.

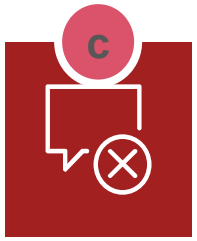

Optimize: The Scaling-Efficiency Paradox

- **Scaling Paradox:**
  - AI workloads are 7-8x more energy-intensive than traditional compute
  - 30x hardware efficiency gains are consistently outpaced by the exponential scale of model deployments.

Source: Data Center Frontier article summarizing IEA report on AI and cryptocurrency doubling data center energy consumption by 2026, International Energy Agency (IEA) report on data center electricity use and AI share of power demand, AI Multiple article "AI Energy Consumption: Statistics from Key Sources [2026]", Pew Research Center article "What we know about energy use at US data centers amid the AI boom", Goldman Sachs Research article "AI to drive 165% increase in data center power demand by 2030", World Resources Institute article "Powering the US Data Center Boom: Why Forecasting Can Be So ...", Electric Power Research Institute (EPRI) paper on future US data center electricity demand, Lawrence Berkeley National Laboratory study on data center demand and U.S. electricity consumption, Congressional Research Service report "Data Centers and Their Energy Consumption" (R48646), European Commission article "In focus: Data centres – an energy-hungry challenge", Google data center efficiency and TPU (Ironwood) energy-efficiency claims, Science Based Targets initiative (SBTi), Integrity Council for the Voluntary Carbon Market (ICVCM), EU Energy Efficiency Directive (recast EED) provisions on data center reporting from 2025; AutoPresent Analysis

The "Offload to Cloud" strategy fails because it merely hides the problem; cloud providers' own emissions are rising, and their carbon intensity is often opaque, preventing true accountability.

	Legacy On-Premise	Hyperscale Cloud
Emissions Ownership	Direct Scope 1 & 2 liability; full visibility into facility efficiency.	Indirect Scope 3 liability; responsibility remains while control is ceded.
Optimization Lever	• Direct hardware tuning • Custom cooling configurations • Physical site selection	• Restricted to API-level settings • Black-box infrastructure • Vendor-defined cooling
Data Fidelity	Real-time, site-specific metering with granular PUE tracking.	• Annualized regional averages • Opaque REC-based accounting • Delayed sustainability reports

30%+

Surge in emissions reported by Google and Microsoft since 2020, fueled by GenAI.

The Transparency Gap

- Carbon intensity is frequently obscured by regional averaging and the purchase of Renewable Energy Credits (RECs).
- Lack of real-time, hourly data prevents CIOs from optimizing compute cycles for low-carbon windows.
- Shifting workloads creates a 'carbon blind spot' where Scope 3 growth offsets Scope 1 reductions.
- Regulators are increasingly focusing on actual physical grid impact rather than paper-based offsets.

The "Offset with Credits" strategy is financially unsustainable, as the cost to offset AI's projected energy growth will eclipse the ROI from the AI initiatives themselves within 5 years.

The "Offset with Credits" strategy is financially unsustainable, as the cost to offset AI's projected energy growth will eclipse the ROI from the AI initiatives themselves within 5 years.

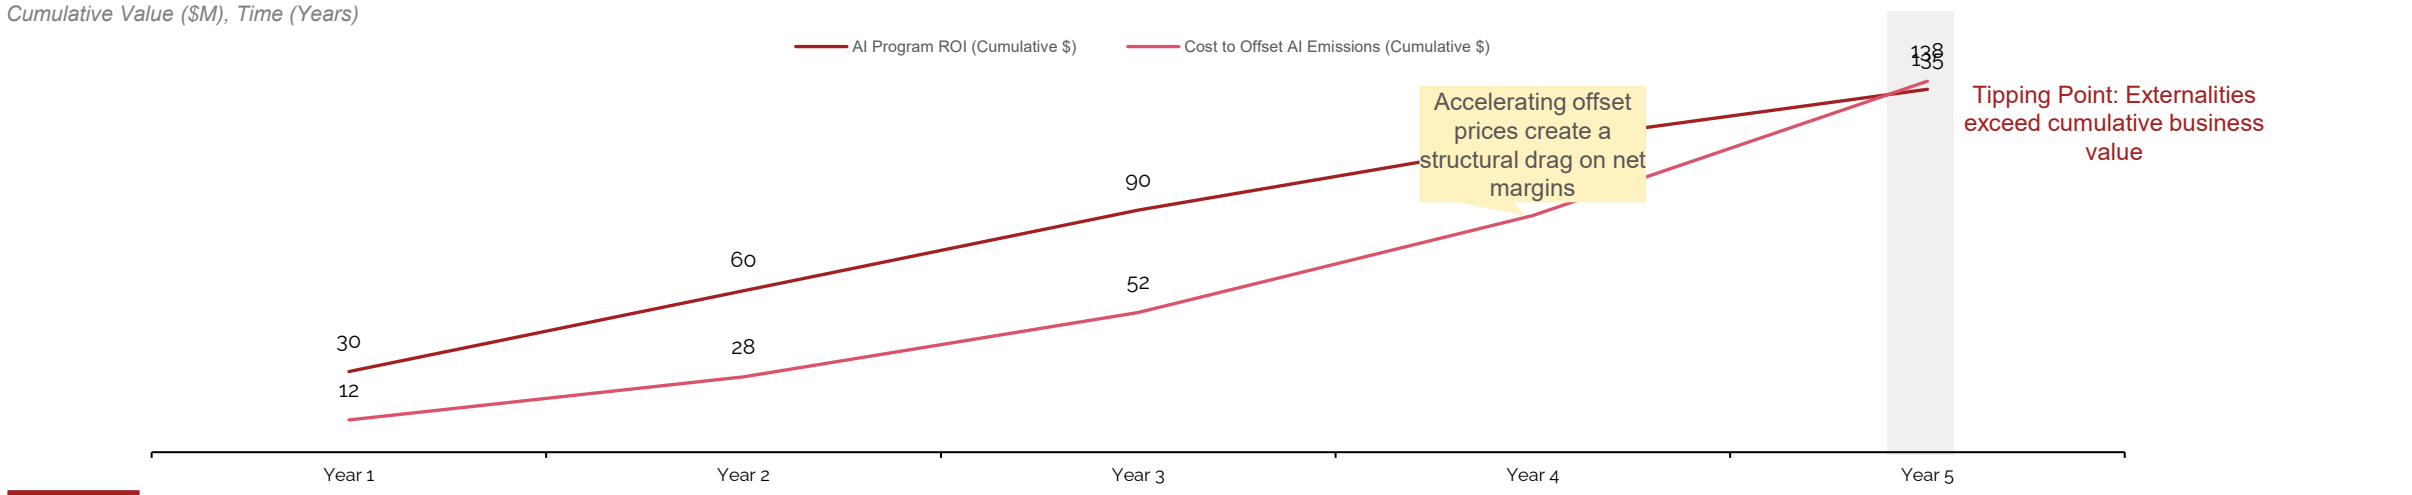

	Year 1	Year 2	Year 3	Year 4	Year 5
Carbon Offset Price ($/tCO2)	30	40	60	90	120
Annual Power Load (TWh)	1	1.1	1.2	1.3	1.4
Net Annual Value ($M)	18	14	6	-11	-33

Negative NPV Threshold: By Year 5, cumulative offset liabilities ($138M) surpass median AI program ROI ($135M), effectively transforming a 15% margin initiative into a net value-destroying asset.

Source: IEA, BloombergNEF, BNEF, Pew, McKinsey, Deloitte, BCG, Carbon Brief, AutoPresent Analysis

The "Optimize in Isolation" strategy is a passive gamble that assumes technology will solve the problem, but breakthroughs in hardware and software efficiency are 5-10 years out, while regulatory pressure is expected within 24 months.

The Strategic Mismatch: Imminent regulatory reporting and grid constraints vs. long-cycle technology diffusion

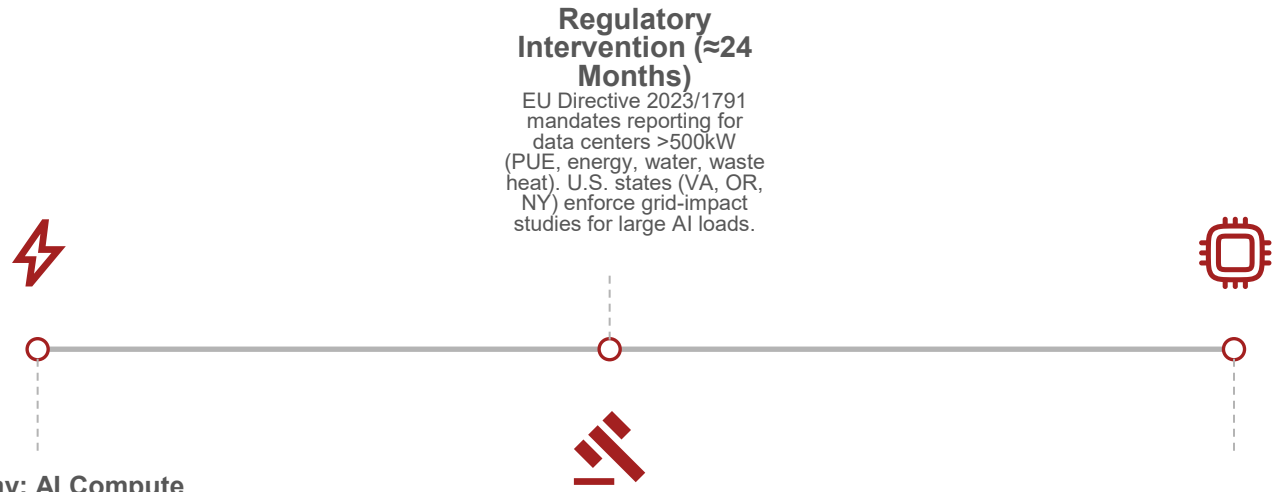

Structural Exposure: Why the 'Passive Gamble' leads to stranded assets and regulatory non-compliance

- **Binding Compliance: EU and U.S. regulatory deadlines arrive within 24 months:**
  - EU Directive 2023/1791 mandates reporting on PUE, water, and waste heat for facilities >500kW by 2027
  - VA, OR, and NY already enforce grid-impact studies, creating de facto growth constraints in key hubs
- **Efficiency breakthroughs remain years away from scale:**
  - Liquid cooling and custom chips (TPUs, NPUs) face hardware refresh and supply chain delays
  - Enterprise-wide impact not expected before 2030 for most
- **AI's energy demand is outpacing all current optimizations:**
  - Data center power use set to double by 2026 despite best-in-class efficiency efforts
  - Incremental gains are being overwhelmed by exponential compute growth
- **The 'Passive Gamble': Deferring structural changes creates existential strategic risk:**
  - Hyperscalers are already revising net-zero strategies as AI loads negate past renewable procurement assumptions
  - Reliance on tech breakthroughs ignores the $7T investment needed in grid and generation infrastructure to meet demand

Source: International Energy Agency (IEA) Report on Data Center Energy Consumption, AI Multiple - AI Energy Consumption: Statistics from Key Sources [2026], Pew Research Center - US data centers' energy use amid the artificial intelligence boom, Goldman Sachs Research - AI to drive 165% increase in data center power demand by 2030, World Resources Institute (WRI) - Powering the US Data Center Boom: Why Forecasting Can Be So Difficult, U.S. Congress - Data Centers and Their Energy Consumption (CRS Product R48646), European Commission - In focus: Data centres – an energy-hungry challenge, Electric Power Research Institute (EPRI) - Data Center Electricity Demand Study, Lawrence Berkeley National Laboratory - Data Center Energy Demand Study, Boston Consulting Group (BCG) - Data Center Energy Consumption Projections, EU Energy Efficiency Directive (EU) 2023/1791, McKinsey Analysis - Technology Diffusion Cycles in Data Center Efficiency, Deloitte - Optimize in Isolation Strategy Analysis, S&P Global - AI Hardware and Software Efficiency Timeline, Stanford Hazy Research - Efficiency Breakthroughs Timeline, AutoPresent Analysis

3

A New Integrated Approach: The Technical Levers and Their Limits

Since all default strategies are flawed, a new integrated approach is required, built on three pillars: Optimize the Compute, Decarbonize the Supply, and Integrate the Governance.

Operationalizing the integrated strategy across infrastructure and oversight silos

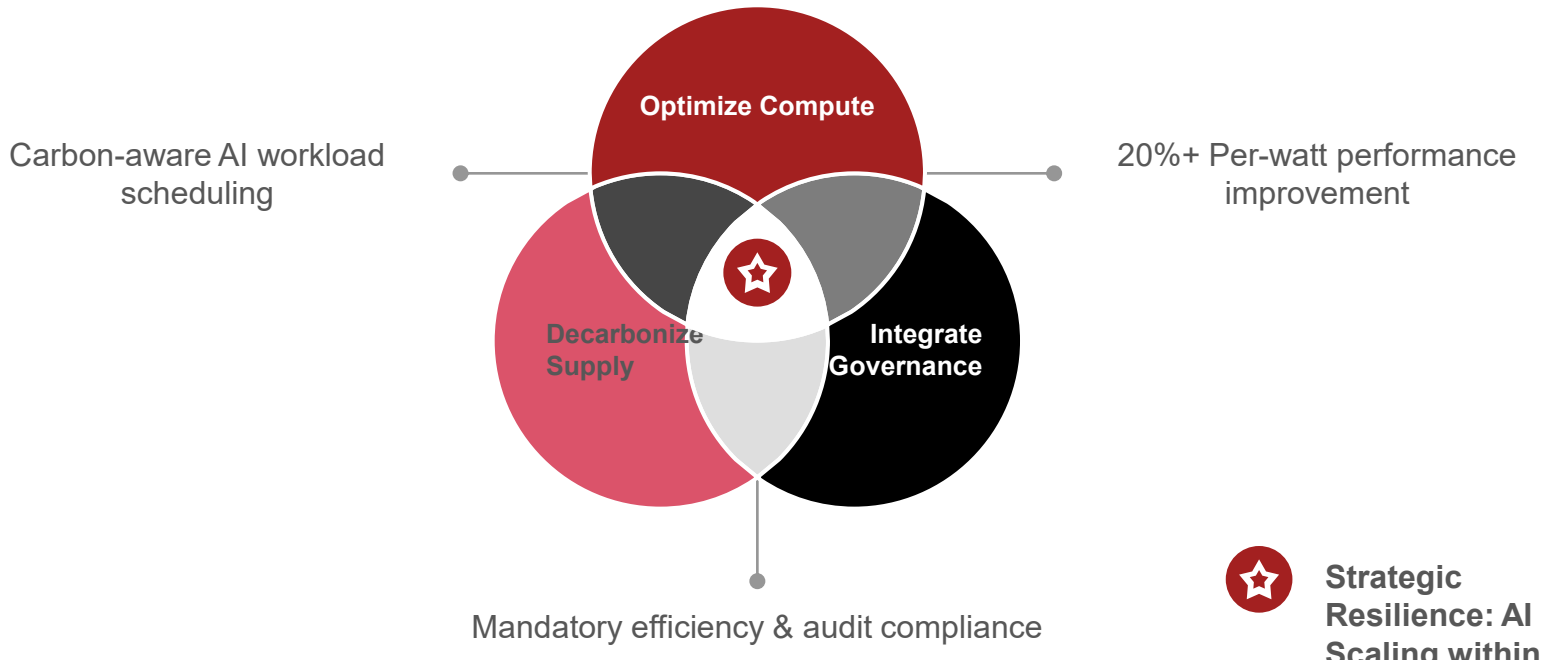

Source: IEA, NVIDIA, Tesla, Thomson Reuters, WEF, HPE's AI-driven Compute Ops Management, NAVER Cloud's sovereign AI, BloombergNEF, Energy Information Administration, Boston Consulting Group, Lawerence Berkeley National Laboratory, AutoPresent Analysis

Rationale: Why integrated governance is the only viable path to scale

- **Default efficiency is outpaced by demand:**
  - Global demand to hit 1,000 TWh by 2026; AI to consume 50% of DC power by 2030
  - 10 GW infrastructure strain requires systemic optimization over localized fixes
- **Supply strategy faces regional grid limits:**
  - Corporate PPAs alone do not guarantee decarbonization due to grid intensity variance
  - 24/7 Carbon-Free Energy (CFE) requires integration with compute demand cycles
- **Regulatory risk mandates new oversight:**
  - Emerging U.S. state audits (CA, VA) and 'critical system' status require cross-functional KPIs
  - Integrated platforms (HPE, NAVER) prove real-time monitoring can align IT with ESG

The first pillar, Optimize the Compute, involves deploying more efficient hardware and smarter software, but each lever has a critical trade-off between performance, cost, and speed of implementation.

Levers offer distinct profiles: Hardware provides foundational efficiency while Software enables tactical agility

Hardware-Led Optimization (Chips & Cooling)	Software-Led Optimization (Quantization & Pruning)
Pros: High impact on watts/PFLOP (~2-3x improvement at chip level) and rack density.	Pros: Material energy savings; stacked techniques (quantization, pruning, distillation) can cut inference energy by 40–80% vs. unoptimized baselines.
Cons: 10-30% up-front CAPEX uplift; system-level gains eroded by networking/cooling overhead.	Cons: Diminishing returns after initial optimization waves.
Implementation Horizon: 24-36+ month lead times for greenfield/liquid-cooled builds.	Implementation Horizon: Near-term wins (weeks to months) utilizing current infrastructure.

Nuance: While chip-level gains reach 2-3x, realized system-level efficiency is significantly lower due to networking and memory overhead.

Grid interconnection upgrades often extend scaling timelines beyond the physical facility build-out, creating a 'hidden' bottleneck.

Strategic Synthesis: Bridge the 2-Year 'Efficiency Gap'

- **Exploit Software Agility Now:**
  - Deploy quantization/pruning to capture 40–80% gains on existing fleets
  - Centralize 'efficiency pods' to overcome talent scarcity bottlenecks
- **Commit to Hardware Foundation:**
  - Initiate liquid cooling builds (24–36m lead time) to lock in 3× floor gains
  - Accept 10–30% CAPEX premiums as prerequisite for high-density AI
- **The 'No Magic Bullet' Mandate:**
  - Parallel tracks required: Software buys time; Hardware ensures survival
  - Treating as an 'either/or' choice creates an unbridgeable power deficit

Source: Data Center Frontier, International Energy Agency (IEA), Lightwave, AI Multiple, Google data centers, Google's latest-generation TPU, Ironwood, Pew Research Center, Electric Power Research Institute, U.S. Energy Information Administration, Goldman Sachs Research, World Resources Institute (WRI), Boston Consulting Group, Lawrence Berkeley National Laboratory, Congressional Research Service report "Data Centers and Their Energy Consumption", European Commission – In focus: Data centres – an energy-hungry challenge, AutoPresent Analysis

However, hardware optimization alone cannot solve the near-term crunch, as the capital expenditure is significant and supply chain lead times for advanced GPUs now extend beyond 24 months.

$200M+
Minimum CAPEX required for a next-gen data center hall.

24+ Months
Supply chain lead times for advanced GPU infrastructure.

Hardware refreshes along cannot solve 2025-2026 emissions gap

- Hardware-centric strategies are structurally incapable of addressing immediate regulatory crises. CIOs must pivot to operational efficiencies and software-level orchestration to mitigate power consumption in the short term while waiting for infrastructure to scale.

By contrast, software optimization is limited by a severe talent shortage and diminishing returns, making it impossible to scale these gains across the entire enterprise AI portfolio.

Software efficiency plateaus at ~45% as marginal engineering effort increases 3–5x

Energy Savings (% reduction vs. baseline), Optimization Effort (Engineering Hours/Complexity)

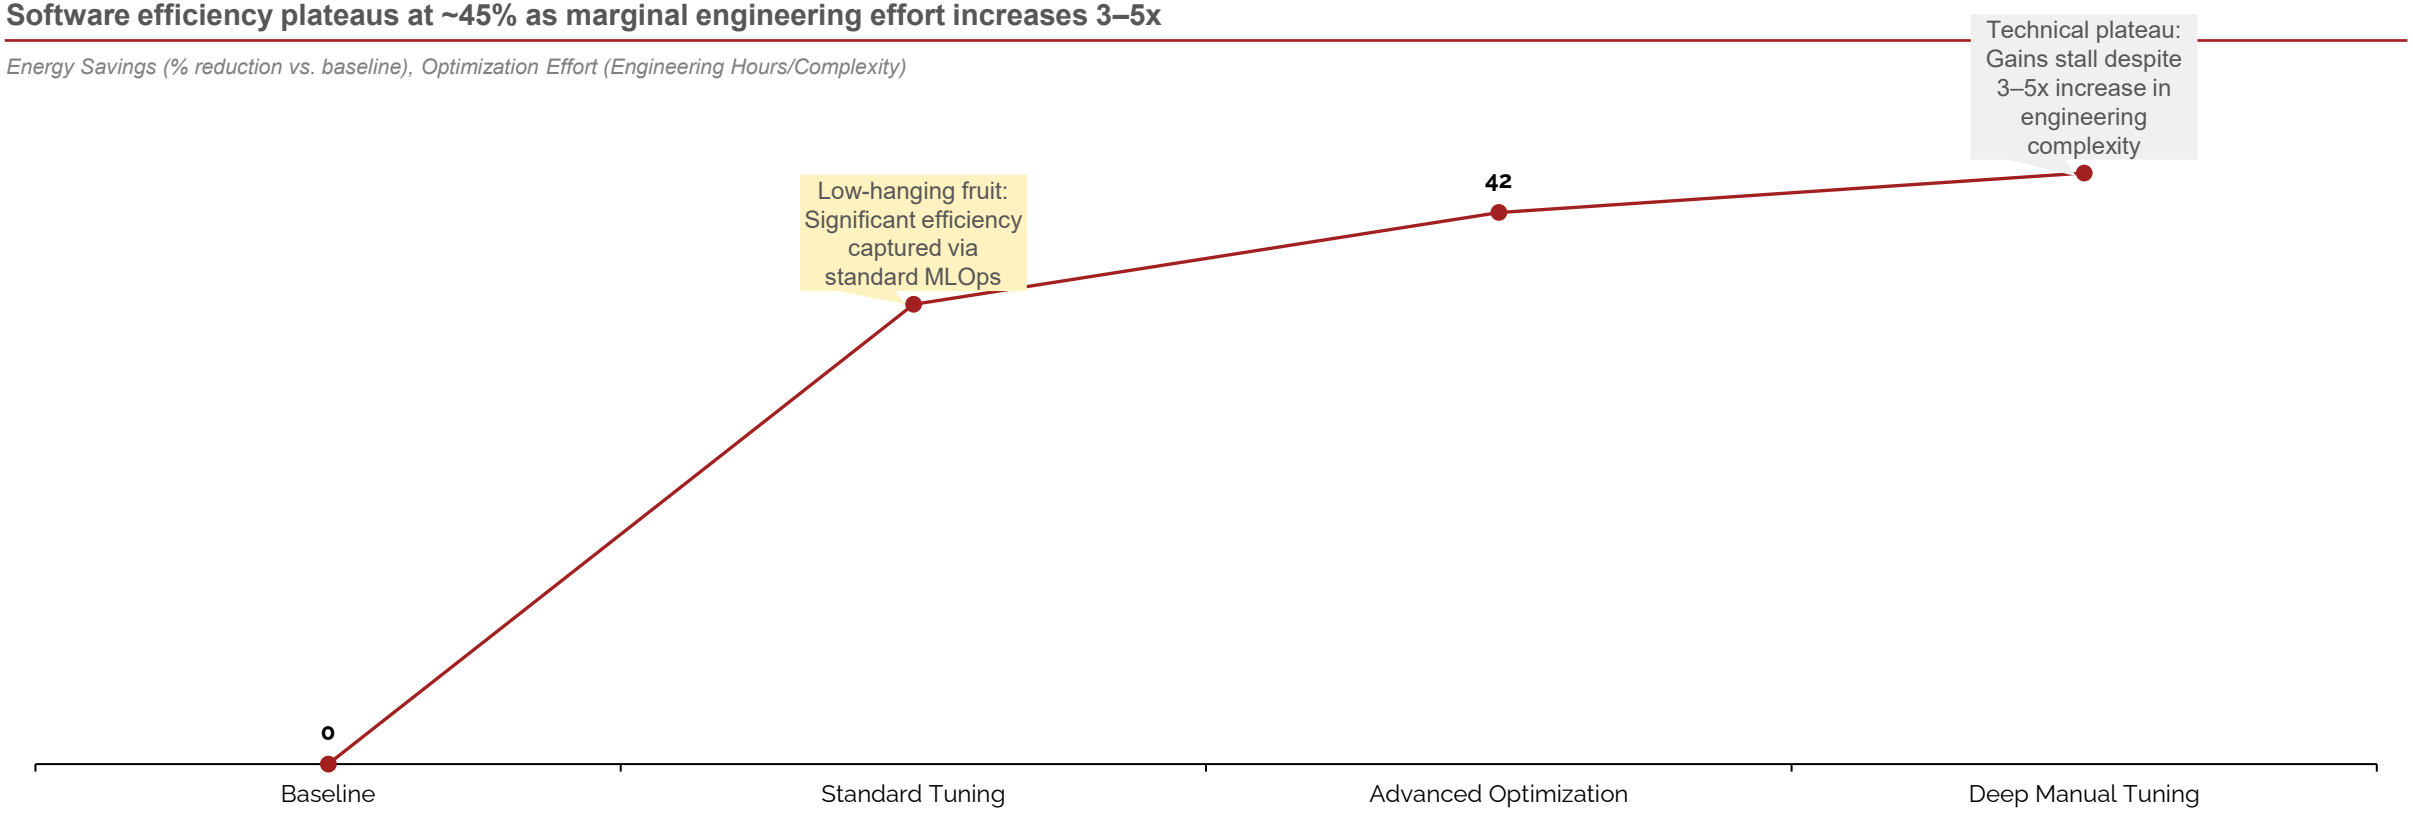

45% ceiling: Software optimization can cut AI energy use by up to 45%—but only if you have unlimited elite engineers; beyond the first 40%, gains stall and the world has fewer than 10,000 people who can push further.

Source: IEEE, ScienceDirect, LinkedIn, Glassdoor, IEEE Access, Elsevier's Sustainable Computing, ACM, TensorRT, XLA, ONNX Runtime, TVM, MobileNet, EfficientNet, VGG, ResNet-50, TensorRT-LLM, vLLM, BetterTransformer, LLaMA-7B, LLaMA-13B, Mistral 7B, AutoPresent Analysis

The second pillar, Decarbonize the Supply, involves securing green power through PPAs and relocating workloads, but this is constrained by PPA scarcity and the inflexibility of existing data center locations.

The Geographic Mismatch: Clean Supply is Physically Remote from AI Load Centers

Strategic Deadlock: Scarcity and Physical Inertia Limit Decarbonization Levers

- **PPA Scarcity: Market demand far exceeds available green supply:**
  - Europe 2024: Corporate PPAs secured only ~20–25% (16 GW) of new wind/solar capacity
  - One AI hyperscale consumes ~1.1 TWh/year—equivalent to the load of 100,000 U.S. households
- **Grid Constraints: Physical bottlenecks create 5+ year delivery delays:**
  - 2,600 GW stuck in U.S. queues; median interconnection wait has surged to 5.2 years
  - Only ~35–45% of N. American corporate PPAs are grid-matched to actual load centers
- **Location Inflexibility: Relocation is a high-cost, high-risk endeavor:**
  - Replacing a 20-MW Tier IV site costs $200–400M in Capex with 60–70% incident risk
  - Unplanned downtime costs exceed $100k/hour; parallel run Opex adds up to $70M

Source: BloombergNEF European Corporate PPA Report (2025), WindEurope – EU wind capacity additions 2024, SolarPower Europe – EU Market Outlook for Solar Power 2024–2028, Lawrence Berkeley National Lab
Agency (IEA) – AI hyperscale data center energy consumption analysis, Uptime Institute – Data center migration cost and downtime risk case studies, Electric Power Research Institute (EPRI) – Virginia data center elect
procurement and grid-matching analysis (North America, 2023), EirGrid / Commission for Regulation of Utilities (CRU) – Ireland data center electricity demand, METI / TEPCO – Japan and Tokyo data center energy use

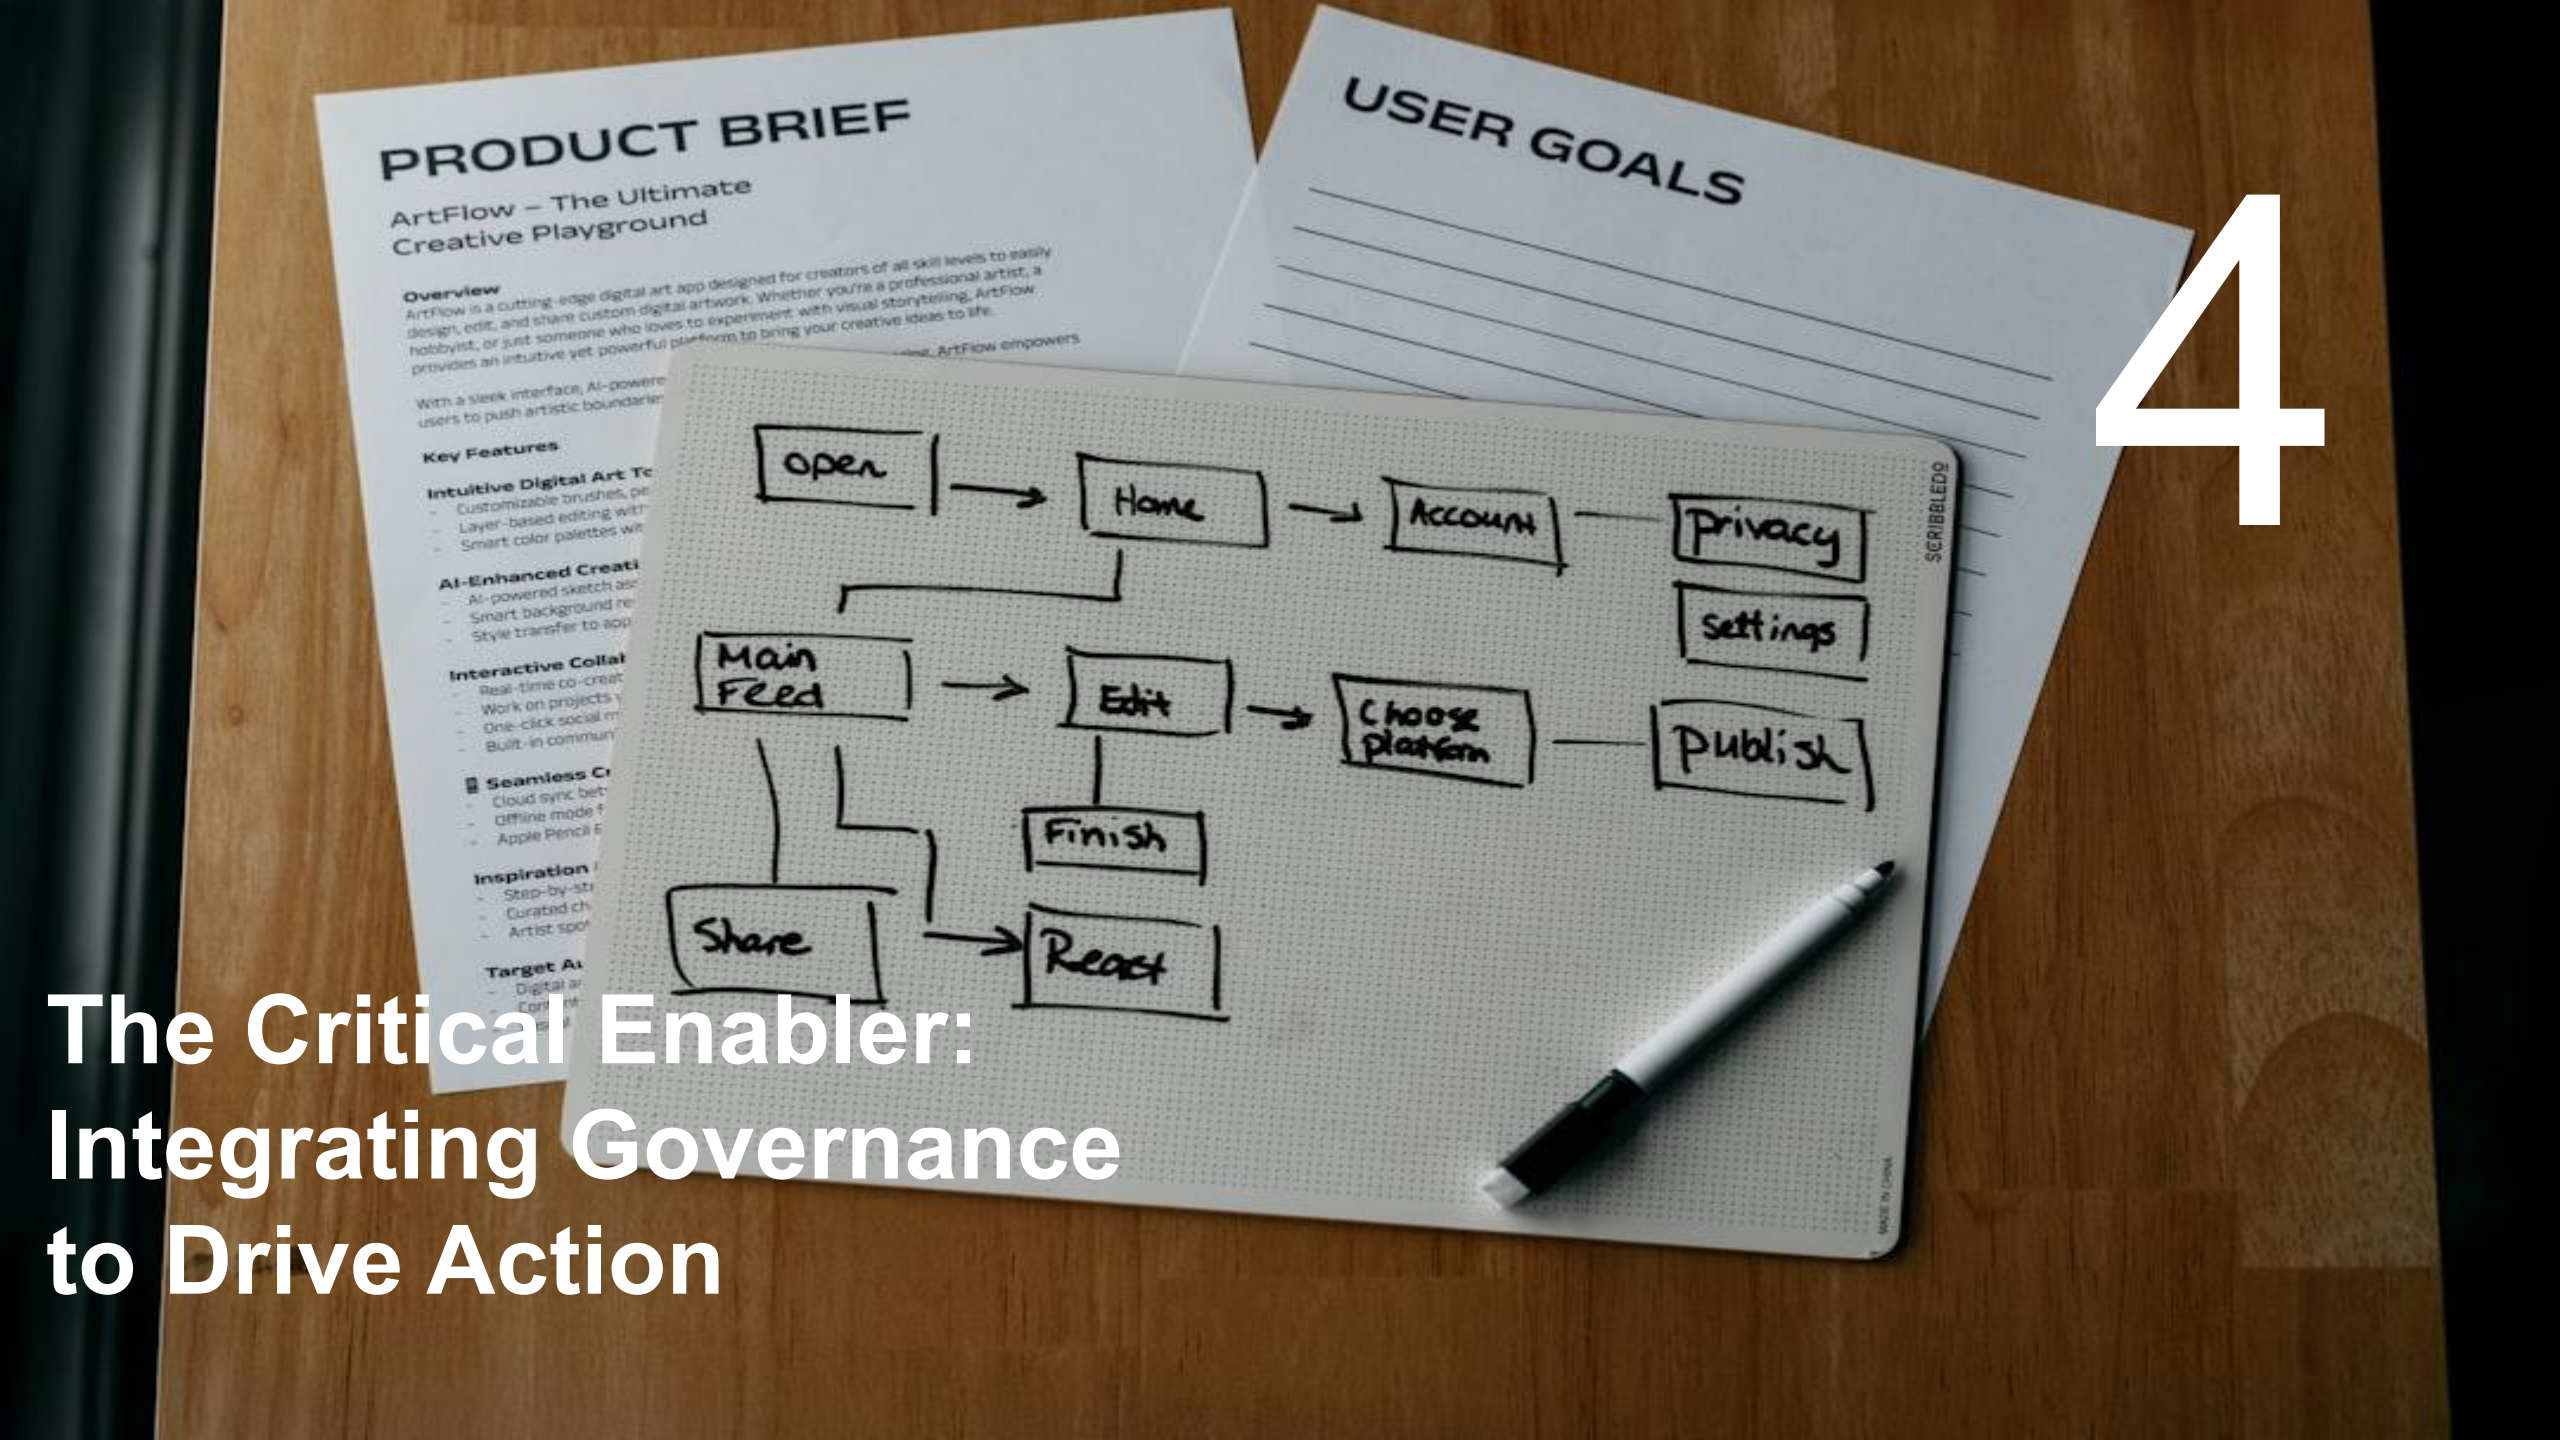

4

The Critical Enabler: Integrating Governance to Drive Action

Because technical levers alone are insufficient, the third pillar, Integrate the Governance, is the critical enabler that forces the necessary trade-offs and makes the first two pillars effective at scale.

The Governance Enabler

01

1. Technical Efficiency (Optimize)

- Demand growth outpaces Moore's Law: US load set to triple to ~580 TWh by 2030, rendering standalone efficiency gains mathematically insufficient.

02

2. Clean Energy (Decarbonize)

- Supply gap: High-density AI scaling (e.g., 7.3 TWh for every 100k servers) overwhelms standard PPA procurement without strategic load triage.

03

3. Integrated Governance (The Enabler)

- Binding Mandate: Unified CIO-CSO decision rights bridge fragmented IT/ESG silos to arbitrate site selection and AI-specific energy KPIs.

3x Demand, 0x Governance:
Without a binding governance model to arbitrate between AI scaling and ESG limits, firms face a 3x power surge that neither PUE gains nor clean energy procurement can bridge alone. Integrated governance is the only lever that forces 'Compute-vs-Carbon' trade-offs at the board level.

Source: IEA, LBNL, NVIDIA, Fortune 500, WSJ 2025, Microsoft, Google, AWS, Publicis Sapient, Liminal.ai, IEEE, internal McKinsey analysis, Electric Power Research Institute, WRI, Boston Consulting Group, U.S. Energy Information Administration, GPT-3, AutoPresent Analysis

This integration is best achieved through a new "Compute & Carbon Board" with joint P&L accountability for AI performance and energy cost, forcing trade-offs to be made explicitly at the executive level.

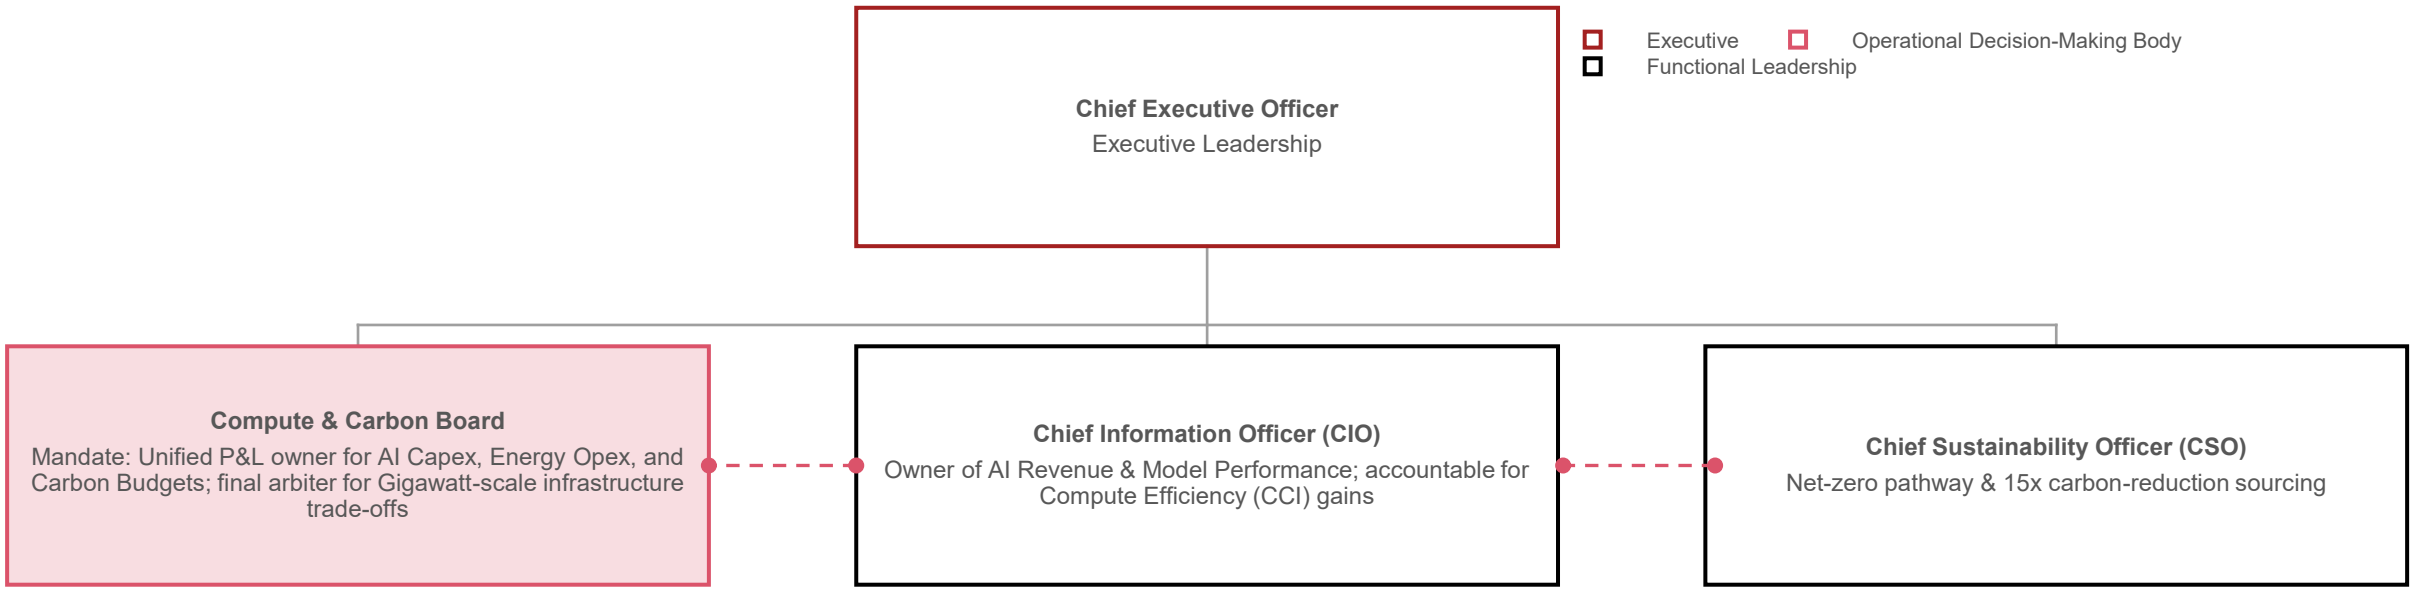

From Advice to Authority:
Moving beyond the partial integration seen at UBS/DBS, this Board closes the 'Agentic AI' governance gap. By unifying 15x carbon-reduction sourcing with AI roadmap speed, it ensures energy scarcity (950+ TWh demand) doesn't stall performance or breach Net-Zero commitments.

Source: IEA Study on Data Center Energy Consumption, AI Energy Consumption: Statistics from Key Sources [2026], US data centers' energy use amid the artificial intelligence boom (Pew Research), Data center energy demand forecasted to soar nearly 300% through 2035 (BloombergNEF Report via TechCrunch), Powering the US Data Center Boom: Why Forecasting Can Be So... (World Resources Institute), Global data center power demand to double by 2030 on AI surge: IEA (S&P Global), Data Centre Energy Use: Critical Review of Models and Results, AI: Five charts that put data-centre energy use – and emissions into context (Carbon Brief), Data Center 2026 Outlook: Energy, Infrastructure, and Connectivity (Morgan Lewis), International Energy Agency (IEA) Reports on Data Center Energy, Electric Power Research Institute (EPRI) Studies, Lawrence Berkeley National Laboratory Data Center Demand Study, Boston Consulting Group (BCG) Data Center Projections, Energy Intelligence (London-based think tank), Google Sustainability Reports and Carbon Compute Intensity Research, Deloitte Analysis on GPU-Augmented Data Centers, AutoPresent Analysis

Precedent for this cross-functional approach exists in leading industrial firms like Schneider Electric, which created a joint business unit for energy and digital, cutting internal energy costs by 25% while growing digital services 15%.

Archetype	Schneider Electric	Unilever
Structural Pivot	Fused energy management and digital architecture (EcoStruxure) into a single P&L	Eliminated standalone CSO; embedded sustainability KPIs directly into CSCO mandate
Efficiency Dividend	25% internal energy cost reduction; 30% consumption drop via AI-load optimization	Decarbonized manufacturing footprint through Scope 1 & 2 operational efficiency
Growth Unlock	15% digital services growth; €6M/year incremental revenue from AI-enabled field services	Secured future growth via 100% sustainable sourcing for key high-risk commodities

Breaking the Trade-off: Integrating energy and digital into a single P&L proves that sustainability is not a growth tax, but a margin lever. This cross-functional model allows firms to fund AI-scale compute expansions through captured energy efficiencies.

Source: Schneider Electric, EcoStruxure, Unilever, S&P 500, TrustCloud, ScienceDirect, AVAIO Digital, AutoPresent Analysis

The CFO's primary objection will be the cost of this new structure, but the business case is clear: avoiding an estimated $50M+ in future carbon taxes and eliminating 15-30% of cloud waste through unified oversight.

	Implementation Cost (The Objection)	Avoided Liability (The Case)
Carbon Tax Exposure	$2M–$5M in annual OPEX for unified monitoring and staffing	$50M+ annual liability (200–300 MW load) as prices reach $80/t (CA) and $125/t (Canada)
Cloud Spend Waste	Minor operational friction from initial FinOps integration	$40M annual saving (20% of $200M spend) via centralized right-sizing
Regulatory Incentives	Internal overhead to manage compliance reporting	• Safeguarding tax breaks in 22 states (60+ bills pending) • Avoided non-compliance penalties

$90M at Stake: Unified governance is a 'Risk Insurance Policy' yielding a ~10x ROI; $90M in annual cost avoidance dwarfs the $5M implementation cost.

Source: CRS report on energy tax benefits for data centers, Energy Tax Benefits for Data Centers: In Brief, Carbon management policy article noting electricity demand projected to rise between 35 and 40 percent by 2040, driven by AI and data centers, What's Next for Carbon Management Policy?, Canada's federal carbon price trajectory to CAD 170/tonne CO_2e by 2030, FinOps industry benchmarks (FinOps Foundation and similar industry surveys on cloud waste 15–30%), Uptime Institute data center PUE benchmarks, Hyperscaler sustainability reports (PUE and grid carbon intensity benchmarks), FinOps Foundation "State of FinOps" reports (cloud cost governance and savings benchmarks), Gartner cloud management and cost optimization reports, IDC cloud spend and market sizing reports, Synergy Research Group public cloud market sizing, Flexera State of the Cloud reports, National Conference of State Legislatures (NCSL) article on state data center incentives and requirements, 2026 Legislative Agendas Put Data Center Incentives in the Spotlight, AutoPresent Analysis

The most significant barrier is cultural; CIOs are rewarded for more compute, CSOs for less, creating a fundamental conflict that can only be resolved by changing executive incentives to a single "Sustainable Performance" metric

The CIO Mandate

- Incentivized for Speed, Innovation, and Uptime.
- Primary KPI: Maximize compute volume and AI model throughput to ensure project delivery and 99.999% system availability.
- Conflict: High energy consumption is viewed as a necessary cost of innovation.

VS

The CSO Mandate

- Incentivized for Reduction, Efficiency, and Compliance.
- Primary KPI: Absolute carbon footprint reduction, resource conservation, and meeting strict ESG regulatory standards.
- Conflict: Unchecked AI scaling is the single largest threat to corporate Net-Zero commitments.

Incentive Reset

- **The Sustainable Performance Metric:**
  - New KPI: (AI Business Value Output / Total Carbon Cost)
  - Synchronizes compute strategy with ESG realities into a single source of truth
- **Redesigning Executive Behavior:**
  - The hardest part is people; shifting from siloed goals to shared accountability for carbon-efficient AI
  - Success must be defined by how much value is created per kilowatt-hour, not just raw uptime

5

Capturing the Advantage: A Roadmap for Sustainable AI Leadership

To have a distinct right to win here, enterprises must use their scale to influence hardware vendors and cloud providers, demanding carbon transparency and more efficient architectures that smaller players cannot.

Scale Advantage	 Purchasing Power	 Cloud Leverage	 Asset Base
Strategic Impact	• Engage in deep co-design with NVIDIA/AMD to set efficiency targets; Google TPU 'Ironwood' delivers 30x energy-efficiency gain over Gen-1 accelerators. • Leverage multi-hundred-billion dollar US hardware capex (2024-30) to institutionalize green computing as a core vendor requirement.	• Move beyond generic dashboards to contractual mandates for per-workload GHG and kWh reporting by region and scope, aligning with OECD/UNESCO standards.	• Deploy 50kW in-rack liquid cooling in owned sites, maintaining PUE of ~1.09 (vs. >30% energy loss to cooling in typical enterprise centers). • De-risk next-gen thermal solutions to drive platform-level transparency and efficiency before industry-wide adoption.

Institutionalizing these requirements creates an efficiency moat: Scaled players lock in lower op-ex and ESG compliance via co-designed hardware and PUEs of ~1.1, while smaller players remain burdened by ~1.5+ PUE legacy sites and standard-tier vendor transparency.

Source: International Energy Agency (IEA) report on data center energy consumption, Goldman Sachs Research article "AI to drive 165% increase in data center power demand by 2030", Google TPU "Ironwood" technical and efficiency disclosures, NVIDIA data center GPU generation comparisons (e.g., V100 vs H100 performance-per-watt), Google Data Centers sustainability and efficiency disclosures (including fleet-wide PUE 1.09), Microsoft Azure data center immersion cooling pilot documentation, Equinix liquid cooling pilot announcements, Google Cloud Carbon Footprint, Microsoft Emissions Impact Dashboard, AWS Customer Carbon Footprint Tool, EU Code of Conduct for Data Centres, OECD AI Principles, UNESCO Recommendation on the Ethics of Artificial Intelligence, Electric Power Research Institute paper on future US data center electricity demand, Lawrence Berkeley National Laboratory study on U.S. data center demand to 2030, EU cloud and data centre public procurement guidance referencing GHG data per service and per location, Supplier codes of conduct and sustainability requirements of Google, Microsoft, and AWS, AI energy consumption meta-analysis "AI Energy Consumption: Statistics from Key Sources", AutoPresent Analysis

To manage concentration risk, this strategy must include a compute portfolio triage, classifying AI workloads as "Critical," "Optimized," or "Frozen" to balance innovation with sustainability goals.

Critical

- **Revenue-generating models:**
  - Focus on core competitive advantages.
  - **Maximize** performance by prioritizing green energy and dedicated high-availability compute to protect top-line growth.

Optimized

- **Internal analytics:**
  - Focus on operational AI and back-office tools.
  - **Implement** aggressive software and hardware efficiency measures (e.g., model pruning) to reduce the total carbon footprint.

Frozen

- **Experimental R&D:**
  - Focus on speculative pilots and non-essential projects.
  - **Pause** compute allocation immediately until the organization's carbon budget or regulatory landscape permits further spend.

This integrated approach also creates adjacent value, enhancing brand reputation with ESG-focused investors and attracting top engineering talent who prioritize working for sustainable leaders.

Sustainability as the Gatekeeper for the $7T Global Data Center Investment Wave

Global ESG Fund AUM (USD Trillion), Year

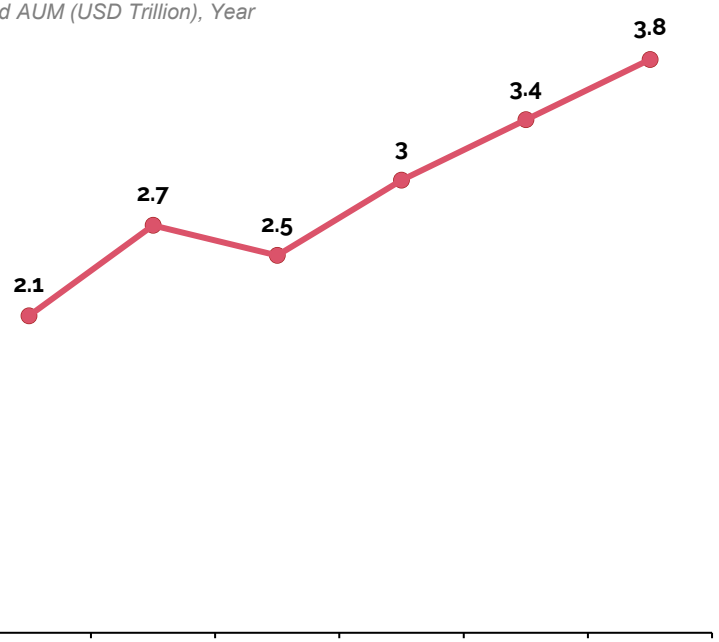

	2020	2021	2022	2023	2024 (E)	2025 (P)
Global DC Power Consumption (TWh)	370	410	460	650	850	980
U.S. Data Center Power Load (TWh)	110	130	150	165	183	220

Lead with Green Compute to Win the War for AI Talent

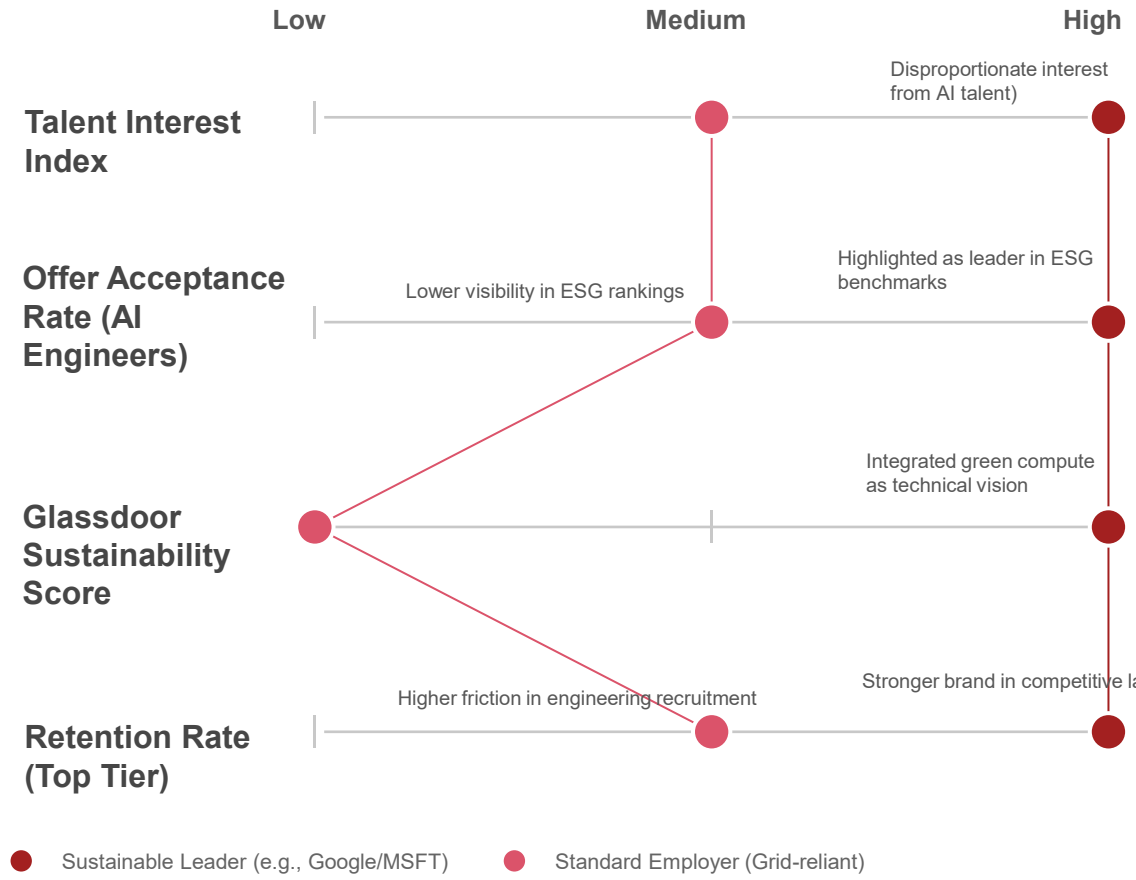

Source: IEA, 2024, U.S. EIA, 2024, McKinsey analysis, 2024, BloombergNEF, 2024, MORNINGSTAR, Bloomberg, EFAMA, ESMA's April 10, 2025 regulatory report on EU sustainable finance, PwC's April 2025 ESG survey, OMRON, Stanford HAI's July 7, 2025 survey on engineering talent priorities, Electric Power Research Institute paper from 2024, Boston Consulting Group, Lawerence Berkeley National Laboratory's study on data center demand, AutoPresent Analysis

The Immediate Action Mandate: The CEO must task the CIO and CSO to co-lead a 60-day sprint to define the charter for the new Compute & Carbon Board, including its unified KPIs and decision rights, for board approval by Q3.

The Sprint Execution Plan: From Mandate to Board Approval

01

Joint Executive Accountability

- CEO tasks CIO and CSO to co-lead a Joint Steering Committee (Model 1), institutionalizing accountability for a projected 165% surge in AI power demand through 2030.

02

60-Day Charter Sprint

- Rapid definition of 'Control Tower' protocols for hyperscale sites consuming ~1 TWh/year—equivalent to 100,000 households—to bridge the current CIO-CSO operational gap.

03

Q3 Board Deliverables

- Finalization of unified KPIs (e.g., gCO2e per AI token) and veto rights for high-compute projects for formal Board adoption by the Q3 meeting.

Strategic Imperative: With data centers set to drive >50% of net electricity demand growth in advanced economies by 2030, this mandate establishes governance before regulatory intervention (CSRD/SEC) or grid constraints force reactive compute cuts.

Source: Goldman Sachs, IEA, LBNL, WRI, CRS, Pew, AI Multiple, SBTi Corporate Net-Zero Standard, CSRD, ISSB, SEC, Lawrence Berkeley National Laboratory, U.S. Congressional Research Service, Pew Research Center article "What we know about energy use at U.S. data centers amid the AI boom", Goldman Sachs Research article "AI to drive 165% increase in data center power demand by 2030", World Resources Institute article "Powering the US Data Center Boom: Why Forecasting Can Be So ...", IEA report on data center energy use (as cited by Data Center Frontier), Google data centers efficiency disclosures (PUE, TPU Ironwood efficiency) as summarized by AI Multiple, AutoPresent Analysis

Questions?