



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 4, April 2025

Carbon Footprint of Digital Infrastructure: Data Centers, Cloud Computing, and the Path to Net Zero

Hitesh Vora

Controls Deployment Engineer, Amazon Data Services Inc, Herndon, VA, USA

ABSTRACT: The exponential growth of digital services—ranging from streaming platforms to business applications—has intensified the energy demand of data centers and cloud computing infrastructures. As governments and private entities commit to net-zero carbon goals, understanding the carbon footprint of digital infrastructure becomes critical. This paper examines the key drivers of energy consumption in data centers, explores current best practices for reducing emissions, and evaluates emerging technological and policy solutions. We conclude that a holistic approach—encompassing energy-efficient hardware, advanced cooling methods, renewable energy sourcing, regulatory incentives, and transparent carbon reporting—can accelerate the path toward net-zero digital infrastructure.

KEYWORDS: Data Centers, Cloud Computing, Digital Infrastructure, Carbon Footprint, Energy Efficiency, Renewable Energy, Net Zero Emissions, Green Computing, LEED Certification, Cooling Technologies, Carbon Usage Effectiveness (CUE), Power Usage Effectiveness (PUE), Sustainability Standards, Environmental Policy, Greenhouse Gas Emissions, Edge Computing, Energy Management, Sustainable Development

I. INTRODUCTION

Over the last two decades, the world has seen an unprecedented surge in data generation. From e-commerce transactions to video streaming services, the modern digital economy runs on cloud-based platforms and massive data centers. These facilities, often spanning hundreds of thousands of square feet, house servers and networking equipment requiring continuous power and cooling. The result is a significant carbon footprint that poses challenges for businesses and policymakers working toward global climate targets.

1.1. The Importance of Digital Infrastructure

Digital infrastructure is integral to every sector—finance, healthcare, education, entertainment, and more. Cloud-based services reduce the need for on-site server installations, thereby offering operational advantages and potential energy-efficiency gains through economies of scale. However, this consolidated demand for computing resources can lead to large, centralized energy loads. As a result, data centers represent an estimated 1–2% of global electricity use, with forecasts suggesting further growth if no mitigation strategies are implemented [1].

1.2. Objective and Scope

This paper aims to:

- Identify the main components contributing to the carbon footprint of data centers and cloud computing.
- Explore current best practices and emerging technologies that reduce carbon emissions.
- Evaluate policy frameworks and corporate strategies aligned with net-zero ambitions.
- Propose an integrated roadmap to guide stakeholders—governments, businesses, and end-users—toward net-zero digital infrastructure.

II. LITERATURE REVIEW

The environmental impact of computing has been a focus of academic and industry research for decades. Early work concentrated on energy consumption within data centers and hardware optimization [2]. In recent years, the scope has broadened to include not only hardware efficiency but also carbon-aware computing and sustainable facility design. Researchers have introduced metrics beyond Power Usage Effectiveness (PUE), such as Carbon Usage Effectiveness (CUE) and Water Usage Effectiveness (WUE), to holistically assess data center sustainability [16]. Studies now address renewable energy sourcing, real-time workload shifting to reduce carbon intensity, and green building certifications that encourage efficient infrastructure [15]. There is also increasing emphasis on “green cloud computing,” which leverages

virtualization and intelligent scheduling algorithms to minimize idle resources [4]. These trends reflect a consensus that both technological innovation and policy frameworks are needed to curb the carbon footprint of digital infrastructure.

2.1. Data Center Energy and Carbon Intensity

Data centers represent a significant and growing portion of electricity demand. Globally, data centers (excluding cryptocurrency mining) consumed on the order of 240–340 TWh in 2022, roughly 1–1.3% of total electricity use ([17]). Broader definitions that include cloud and high-performance computing workloads estimate usage closer to 2% of global electricity (about 460 TWh in 2022) [8]. In the United States, data center energy consumption has surged in recent years. After plateauing at ~60 TWh in 2014–2016, U.S. demand grew to approximately 176 TWh by 2023 ([17]), about 4.4% of U.S. electricity use. [7] This rapid growth is attributed to the expansion of hyperscale cloud facilities and energy-intensive computing tasks. Projections are sobering: U.S. data centers may reach 6–12% of national electricity use by 2028 if current trends continue [3]. Globally, if efficiency improvements stagnate, data center demand (together with emerging AI workloads) could approach 3–4% of electricity use by the end of the decade [8]. The carbon emissions associated with this energy use are substantial. Data centers and related digital infrastructure accounted for roughly 330 Mt CO₂ in 2020 [16], on the order of 1% of energy-related greenhouse gas emissions. By 2022, combined emissions from data centers and crypto operations were nearly 1% of the global total [11]. If no mitigating actions are taken, data center CO₂ emissions could reach 450 Mt by 2027 (~1.2% of world emissions) [11]. These figures underscore the importance of improving energy efficiency and decarbonizing the power supply for data centers.

What is the average annual PUE for your data center? (n=567)




Figure 1. Industry-wide PUE improvement slows, then stops [19]

Industry-wide average Power Usage Effectiveness (PUE) over time. The global **average PUE** dropped from ~2.5 in 2007 to ~1.65 by 2014 as best practices in cooling and power management were adopted, but has since plateaued around 1.55–1.6 in recent years [19]. This stagnation suggests that further efficiency gains must come from IT equipment improvements and renewable energy, rather than cooling/power alone.

Despite the plateau in PUE, there are signs of progress in select areas. Hyper-scale cloud operators achieve significantly lower PUE values by leveraging custom engineering and scale. For instance, Meta's data centers (which are all LEED Gold or higher certified) operated at an average PUE of just **1.09** in 2023 ([14]). This indicates state-of-the-art facilities can approach the theoretical efficiency limit ($PUE \approx 1.0$) where nearly all energy powers IT equipment. However, industry surveys confirm that many legacy or smaller data centers still run with PUE in the 1.7–2.0 range, especially in regions with hot/humid climates or less efficient cooling [19]. As a result, the global fleet's efficiency gains have been offset by the expansion into regions and workloads that are less optimized. Recent literature calls for new strategies to break through the stagnation, including improving server utilization (reducing energy per computation) and incorporating clean energy to cut carbon per kWh [16]. Researchers also note the importance of complementary metrics like CUE (which tracks kgCO₂ per kWh consumed) to ensure that a low PUE actually translates into low carbon footprint [16]. In summary, while energy efficiency improvements have moderated the growth of data center energy use, continued efforts are required to decouple digital infrastructure expansion from carbon emissions.

2.2. Cooling and Thermal Management

Significant attention has been devoted to the cooling infrastructure, as cooling can account for 30–40% of a data center's energy consumption. Traditional air-cooling systems (CRAC units, chillers, air handlers) are energy-intensive, particularly in large facilities or warm climates. Innovations in this domain have been a major focus of recent research. Advanced **liquid cooling** and **immersion cooling** techniques can dramatically improve heat removal efficiency, allowing higher rack power densities with less energy input. Studies show that immersion cooling (submerging servers in dielectric fluid) can reduce cooling energy usage by 30–50% . [12] "Free cooling" methods are increasingly adopted—utilizing outside air or cold water sources to dissipate heat without compressor-based chillers. For example, some Nordic data centers use winter air or sea water to cool servers almost year-round, yielding PUE improvements and lower cooling costs [20]. Large cloud operators often design facilities to take advantage of ambient conditions; **Meta's** data centers, for instance, use indirect evaporative cooling with outside air, enabling them to use **80% less water** than typical data centers while maintaining efficient cooling [21].

Another emerging approach is locating data centers in naturally cold regions or underground. Research on **geo-temperature cooling**, including projects like Microsoft's submerged ocean data center experiment, suggests that unconventional siting can nearly eliminate active cooling needs, though such approaches are in early stages. Even in conventional sites, improved airflow management (hot/cold aisle containment, raised floors, blanking panels) has matured as a best practice to minimize mixing of hot and cold air streams [19]. This, combined with smarter cooling controls (dynamic fan speeds, AI-optimized cooling setpoints), has yielded incremental gains. The net effect is that modern facilities can often keep cooling PUE (cooling energy/IT energy) well below 0.5, whereas older designs might have cooling overhead approaching parity with IT load. Going forward, cooling research is focusing on the needs of high-density racks for AI and HPC: liquid cooling of chips, two-phase immersion cooling, and even on-chip microfluidic cooling to handle heat fluxes that air can no longer practically remove [20]. Efficient thermal management remains critical, as every watt saved in cooling reduces total energy use and indirectly cuts carbon emissions (particularly if the grid power has a carbon factor).

2.3. Policy and Regulatory Frameworks

Governments worldwide are developing policies and incentives to curb data center emissions. In Europe, directives now require large data centers to measure and report their energy use and carbon emissions. For example, the European Union's latest Energy Efficiency Directive mandates that data centers over 500 kW report their energy consumption and waste-heat reuse metrics, with compliance starting in 2024 [8]. Such transparency requirements aim to foster competition on efficiency and enable policy makers to track industry progress. Some regions are also enforcing **carbon caps** or integrating data centers into emission trading schemes. Singapore famously imposed a moratorium on new data centers in 2019 due to energy concerns and now only approves new projects that meet strict efficiency criteria (e.g., PUE targets and renewable energy use) [15]. The **Climate Neutral Data Centre Pact** in Europe (a voluntary agreement) commits operators to ambitious efficiency and green energy goals by 2030 effectively self-regulating in anticipation of stricter laws.

In the United States, policy approaches have included both incentives and possible regulations. Some U.S. states offer tax breaks for data center investments contingent on sustainable design features or renewable energy procurement agreements. At the federal level, the Department of Energy has highlighted data centers in its energy strategy, recently reporting that data center load growth could significantly impact national power demand by 2028 [18]. While no federal mandates on data center emissions exist yet, the **ENERGY STAR for Data Centers** program provides an efficiency certification for facilities, and there are discussions about requiring large data centers to conduct energy audits or adhere to building energy codes. The interplay of regulation, market forces, and consumer demand for greener cloud services is clearly pushing operators toward more sustainable practices [9]. Importantly, new standards and frameworks are not only targeting energy use but also water usage and e-waste. For instance, some jurisdictions are considering water usage reporting requirements for data centers, reflecting concerns over the **water footprint** of cooling systems. In summary, policy frameworks are evolving rapidly: they range from soft measures (reporting and voluntary standards) to hard measures (caps and efficiency mandates), all geared toward aligning the digital infrastructure sector with climate goals.

III. METHODOLOGY

This paper employs a multi-pronged research approach:

1. **Literature Synthesis:** We reviewed peer-reviewed journals, industry reports, and policy documents to identify the key drivers of carbon footprints in digital infrastructure. Special attention was given to studies published in the last five years to incorporate the latest data on energy usage trends and emerging mitigation technologies.
2. **Case Studies:** We examined select data centers that have implemented best practices in energy efficiency, innovative cooling, and renewable integration. These case studies (e.g., a hyperscale cloud data center vs. a smaller edge data center) illustrate the range of current performance and the impact of specific interventions.
3. **Comparative Analysis:** We assessed various carbon-reduction strategies—both technological and policy-driven—using metrics like energy intensity (kWh per compute workload), carbon savings (Mt CO₂ avoided), and cost-effectiveness. This involved comparing scenarios (business-as-usual vs. green-optimized) using data from the literature and industry benchmarks.
4. **Roadmap Development:** We synthesized the findings to propose an integrated framework guiding stakeholders (governments, businesses, and end-users) toward net-zero digital operations. This roadmap considers near-term and long-term actions, and it aligns technological solutions with policy measures and corporate sustainability commitments.

This mixed-method approach enables both a quantitative understanding of energy and carbon trends and a qualitative insight into the strategies that work in practice. By combining data analysis with real-world examples, the methodology ensures that conclusions are grounded in evidence and applicable to industry and policy decision-making.

IV. KEY DRIVERS OF CARBON FOOTPRINT IN DIGITAL INFRASTRUCTURE

A data center's carbon footprint is determined by a combination of its energy consumption and the carbon intensity of that energy. Several core components and operational factors drive this footprint:

4.1. Server and Networking Equipment

Servers are the heart of data centers, performing computations from basic web hosting to complex AI model training. As compute demands rise, the number and performance of servers increase, leading to higher energy draw. Even idle or underutilized servers consume significant power (often 30–50% of peak power) due to baseline motherboard, disk, and power supply loads. This means that **low server utilization** can greatly inflate a facility's energy per useful compute. Networking equipment—switches, routers, storage arrays—adds to the power draw, and their efficiency varies. Legacy hardware tends to be less energy-efficient; thus older data centers with outdated servers and network gear have a higher watts-per-workload than modernized facilities. A key driver here is the **compute efficiency** of hardware. Recent improvements, such as more efficient CPU architectures and purpose-built accelerators, can reduce the energy required for the same amount of processing. For example, cloud operators have introduced ARM-based server processors that use up to 60% less energy for equivalent performance compared to traditional x86 servers [22]. Such gains directly lower the carbon footprint per computation when coupled with high utilization via virtualization. On the other hand, **demand for high-performance computing (HPC)** and low-latency networking can push power use upward. In summary, the balance of modern, energy-efficient hardware and how well that hardware is utilized is a primary driver of a data center's carbon profile.

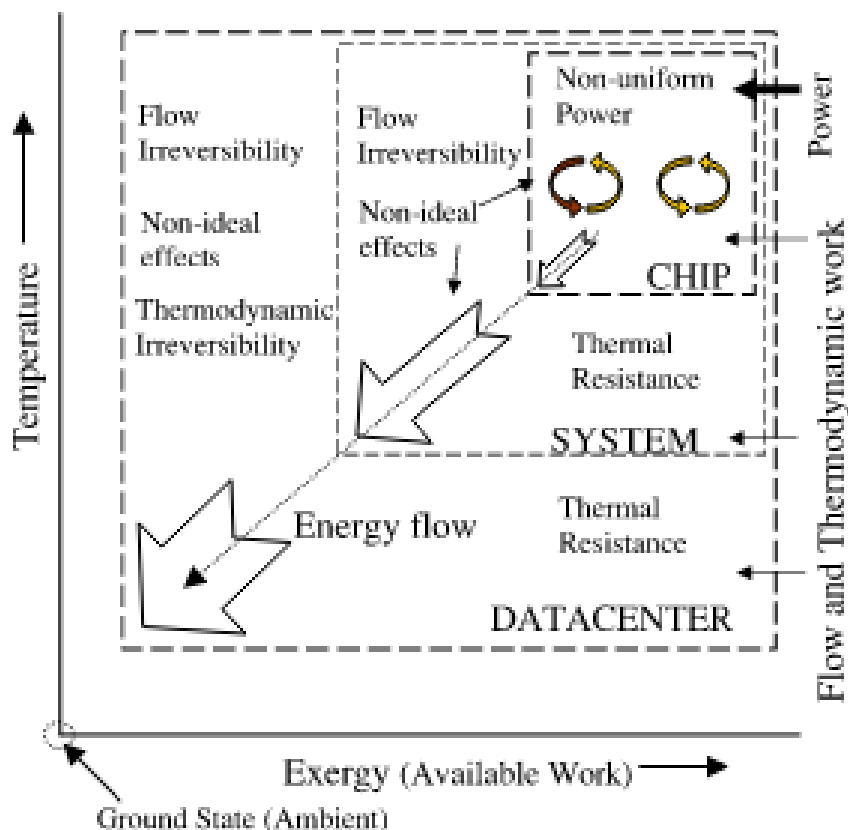


Figure 2: Schematic of exergy destruction and energy flow from chips to data center [6]

4.2. Cooling Systems

Maintaining servers at safe operating temperatures is essential for reliability, and it has a major influence on total energy use. Traditional data centers rely on chillers, Computer Room Air Conditioning (CRAC) units, and air handlers that can consume almost as much energy as the IT equipment if not optimized. An inefficient cooling system (e.g., poor airflow management, over-cooling to unnecessarily low temperatures) drives up energy use and thus carbon emissions if the power is from fossil fuels. The **Power Usage Effectiveness (PUE)** metric is heavily influenced by cooling efficiency – a high PUE often indicates excessive cooling overhead. Elevated-temperature designs (allowing server inlet temps of 27 °C or higher per ASHRAE guidelines) and advanced cooling methods significantly reduce this overhead [20]. For instance, using liquid cooling or rear-door heat exchangers can remove heat more directly and with less fan power than traditional room air cooling. Free cooling (direct use of outside air or water) can, in suitable climates, eliminate chiller use for large portions of the year [20]. The **carbon impact** of cooling also depends on refrigerants and backup systems; leaks of refrigerants (which are potent greenhouse gases) and inefficient backup chillers can add to the footprint. Thus, the design and operation of the cooling system—airflow layout, setpoints, equipment efficiency, use of economization—are critical drivers of both energy consumption and carbon emissions in data centers.

4.3. Power Distribution and Losses

Power must be delivered from the grid (or on-site generation) to every server, and inefficiencies in this chain contribute to the carbon footprint. In large facilities, power distribution involves transformers, switchgear, power distribution units (PDUs), and voltage conversion steps (e.g., from high-voltage AC to server-level DC). Each step can introduce losses (heat), often on the order of a few percent. While modern electrical designs achieve high efficiency (over 95% at each conversion), older or poorly loaded infrastructure may incur greater losses. **Uninterruptible Power Supplies (UPS)**, which provide battery backup, also impact efficiency – double-conversion UPS systems typically have 8–10% loss, though newer high-efficiency or “eco-mode” UPS can reduce this. Redundancy (N+1 or 2N power paths) is crucial for reliability but means some equipment runs at partial load (where it’s less efficient) and duplicates energy losses. All these

electrical losses add up as heat, which then further loads the cooling system. Thus, a data center with inefficient power distribution has a twofold carbon impact: direct energy waste and increased cooling requirements. Designing power delivery with higher efficiency components (e.g., high-efficiency UPS modules, minimizing AC-DC conversion stages, using higher-voltage distribution to reduce resistive losses) can reduce these losses. For example, moving from a legacy UPS (90% efficient) to a modern UPS (>97% efficient) can save millions of kWh annually in a large data center [19]. In sum, power distribution losses, while a smaller fraction of total energy use than IT or cooling, are a notable driver of non-IT energy consumption and thus contribute to the overall carbon footprint.

4.4. Backup Generators and Emissions

Most data centers rely on backup diesel generators to ensure continuous operation during grid outages. While generators are infrequently used (often just tested periodically and during rare emergencies), they contribute to the facility's potential carbon emissions and local air pollution. When they do run, diesel gensets emit CO₂, particulate matter, and NO_x. A single large data center can have tens of diesel generators, which collectively emit significant greenhouse gases if operated for extended periods (for instance, during prolonged grid failures or for load management). In terms of the overall carbon footprint, generator use is usually a **minor contributor** (since annual runtime is low) but remains a point of concern for **carbon accounting and air quality compliance**. Some jurisdictions restrict generator runtime due to air pollution, which in turn pressures data centers to find cleaner backup solutions. Operators are now experimenting with alternatives like **hydrogen fuel cells** and large-scale battery storage to replace diesel units. These technologies promise zero on-site emissions (hydrogen fuel cells emit only water, and batteries emit nothing on-site) and could drastically reduce the carbon footprint associated with backup power. However, until such alternatives are widely deployed, diesel generators remain an emissions driver—especially if a data center participates in demand response programs that deliberately switch to generators during peak grid periods (effectively shifting load off the grid but burning diesel). Therefore, backup power systems and their fuel sources are part of the carbon footprint equation, albeit usually latent. Transitioning to cleaner backup power not only cuts potential carbon emissions but also aligns data centers with air quality regulations and corporate sustainability goals.

4.5. High-Density Compute Workloads (AI and HPC)

An emerging driver of digital infrastructure's carbon footprint is the rise of high-density compute workloads, such as artificial intelligence (AI) training, machine learning, and cryptocurrency mining. These workloads typically run on power-hungry hardware like GPUs and specialized accelerators that can draw significantly more power per server than traditional enterprise workloads. For example, a single AI training server with multiple GPUs can consume 3–5 kW or more, an order of magnitude above a typical 1U server from a decade ago. The growing demand for AI services (e.g., large language model training and inference) has led cloud providers to deploy entire clusters of such high-power servers, driving up the energy intensity of data centers [8]. Studies have highlighted how **AI workloads can dramatically increase energy consumption** – one estimate noted that an advanced AI query (ChatGPT-level) may use 10× the electricity of a standard web search [11]. The carbon footprint of these workloads is magnified if the additional energy is drawn from fossil-fueled grids. Cryptocurrency mining, while not a traditional “cloud” workload, also contributes heavily; globally, crypto data centers (mining farms) consumed on the order of 100–150 TWh in 2022 [17], often with a high carbon intensity if powered by coal-dominated grids (e.g., some Bitcoin mining operations).

Thus, the **compute density** (watts per server and watts per rack) is becoming a key driver of infrastructure energy use. A data center full of AI servers will have a far larger power draw (and potentially carbon footprint) than one of similar size running low-intensity workloads, unless mitigated by sourcing carbon-free energy. This trend pushes operators to adopt advanced cooling (to handle the heat from dense racks) and to double down on renewable energy procurement (to offset the significant carbon emissions from the power required). In summary, the changing mix of workloads—toward more computationally intensive tasks—is a modern driver of increased energy use and carbon emissions in data centers, one that must be addressed through both technology (efficient hardware, better algorithms) and energy strategy (ensuring these new loads are met with green power).

V. STRATEGIES FOR REDUCING DATA CENTER EMISSIONS

Addressing the carbon footprint of data centers requires a multifaceted approach. Key strategies include technological innovations to improve efficiency, transitioning to clean energy sources, and intelligent management of IT workloads and cooling. Below we evaluate major strategies, noting recent developments and their effectiveness in reducing energy use and emissions:

5.1. Energy-Efficient Hardware and Virtualization

Adopting energy-efficient servers and networking gear is a fundamental step to reduce baseline power draw. Modern CPUs and GPUs have improved performance-per-watt characteristics, partly through smaller semiconductor process nodes and architecture enhancements. Many incorporate dynamic power management features (dynamic voltage/frequency scaling, core parking) that allow them to throttle down during low utilization. This reduces wasted energy when full performance is not needed. Additionally, chip manufacturers and cloud companies have developed specialized processors (such as AWS's **Graviton** ARM-based CPUs and Google's Tensor Processing Units) that are tailored for high efficiency on certain workloads. For example, AWS reports that its Graviton3 processor uses **60% less energy** for the same performance compared to a traditional EC2 instance [22], significantly cutting the energy required per computation. On the virtualization side, techniques that consolidate workloads are crucial. Virtual machines and containers enable multiple applications to run on a single physical server, increasing average utilization.

Higher server utilization means fewer total servers are needed for a given workload, directly reducing energy use. Cloud providers routinely run physical hosts at much higher utilization (40–70%) than typical on-premises enterprise data centers (which might average 10–20% utilization) [23]. This consolidation, powered by robust virtualization and container orchestration, translates to efficiency gains. Another aspect is **decommissioning or upgrading inefficient hardware**: replacing older servers (which may consume the same power as new ones but deliver less performance) can improve overall compute efficiency. In practice, leading operators like Meta have reported that their data center refreshes and custom server designs make their facilities 22% more energy-efficient than the industry average [14]. Ultimately, by pairing efficient hardware with aggressive virtualization and cloud computing models, data centers can perform the same amount of computational work with a fraction of the energy (and associated emissions) that would be required in a less optimized environment.

5.2. Renewable Energy Procurement

Switching the power supply of data centers to renewable energy is one of the most impactful ways to cut carbon emissions. Because even an efficient data center draws a lot of power, ensuring that power is generated from carbon-free sources (solar, wind, hydro, etc.) can zero out the operational carbon footprint (Scope 2 emissions). Data center operators are increasingly procuring energy from renewables to meet sustainability targets. Major cloud providers have become **leading corporate purchasers of renewable energy** [25]. For instance, Amazon, Microsoft, and Meta were the top three buyers of wind and solar power purchase agreements (PPAs) in 2021 [25] collectively securing gigawatts of green power capacity. These long-term contracts finance new renewable projects, bringing additional clean energy onto the grid. Google has matched 100% of its annual electricity consumption with renewable energy since 2017, primarily through PPAs, and is now striving for 24/7 carbon-free energy by 2030 (ensuring every hour of consumption is met with local clean generation) [10]. In practice, renewable procurement can take several forms: utility-scale PPAs, on-site generation (e.g., rooftop solar or fuel cells), or purchasing renewable energy certificates (RECs) to offset grid power use. Many data center companies are also investing directly in solar and wind farms. The scale is notable – the top five hyperscalers (Amazon, Google, Meta, Microsoft, Apple) had a combined renewable portfolio of over 45 GW by 2022 [26] an amount rivaling utilities. For smaller operators or colocation providers, joining aggregated procurement groups or buying green tariffs from utilities are common strategies.

The key is **additionality**: ensuring the renewable energy is actually adding new clean generation rather than simply diverting existing capacity. When done properly, renewable sourcing can drastically reduce emissions; a data center running on 100% renewable electricity has effectively zero Scope 2 carbon emissions. However, challenges remain in achieving round-the-clock renewable supply (hence the push for energy storage and flexible load shifting). Some operators also explore **on-site energy storage** so they can ride through intermittency and even use batteries to shave peak grid demand (further supporting renewable integration). Overall, renewable procurement has become a cornerstone of corporate climate strategies in the data center industry, evidenced by initiatives like Google's carbon-intelligent computing and Microsoft's commitment to be carbon negative (removing more CO₂ than it emits) by 2030 [13]. The success of these efforts is reflected in metrics like CUE (carbon usage effectiveness) improvements and in the public reporting: for example, Google in 2022 achieved a 64% carbon-free energy rate across its data centers [27], and Meta and others report 100% renewable energy matching for their operations [14].

5.3. AI-Driven Workload Management

Cloud computing offers the flexibility to distribute tasks across many servers and even multiple data center locations. Leveraging this flexibility, AI-driven workload management has emerged as a powerful technique to optimize energy

usage and reduce emissions. The idea is to use machine learning algorithms and predictive analytics to dynamically schedule and place workloads in a way that either improves energy efficiency or shifts compute to times/places with cleaner power. One approach is often called “**follow-the-sun**” or “**follow-the-renewables**”, where non-urgent computing tasks are scheduled when renewable energy is abundant (e.g., midday when solar is at peak, or at night if there’s excess wind power) [10]. Google implemented a carbon-intelligent load shifting system that delays certain batch workloads until the grid carbon intensity is lower [10]. Early results showed this strategy can significantly cut emissions without affecting service levels [10]. AI can also predict demand patterns and pre-emptively consolidate workloads during off-peak hours, allowing some servers to enter low-power idle states or be turned off completely. Furthermore, across geographic regions, AI systems can route new requests to data centers that currently have spare capacity or lower marginal carbon emissions. For example, if a user’s request can be served from either of two cloud regions, the system might choose the region where it is nighttime (cooler, less air conditioning needed) or where the grid is greener at that moment.

This type of **geographical load balancing** has been shown to reduce overall carbon footprint when implemented intelligently [10]. AI-driven management is also used within the facility for **cooling optimization** – so-called “smart cooling” that adjusts setpoints and cooling unit output based on real-time IT load and weather forecasts, something that Google and others have deployed using DeepMind AI, yielding energy savings of 30% in cooling systems. In essence, AI provides a granular control layer that can squeeze out inefficiencies that static or manual operations would miss. By continually learning and adjusting, such systems ensure data centers run as leanly as possible. The combined effect of AI-driven workload and cooling management can be a substantial reduction in energy use (thus emissions) – one study found that carbon-aware scheduling could reduce a data center’s emissions by 10–20% without additional hardware [10]. As demand for AI services grows, ironically AI itself will be a key tool in managing the sustainability of the underlying infrastructure.

5.4. Advanced Cooling Techniques

Cutting-edge cooling technologies offer significant opportunities to reduce data center energy use and carbon impact, as mentioned in Section 2.2. Here we highlight some advanced techniques and their benefits:

1. **Liquid Cooling:** By bringing coolant directly to the heat source, liquid cooling removes heat more efficiently than air. Direct-to-chip cold plates or immersive baths can keep server components at lower temperatures with far less airflow. Studies have quantified up to *30–50% reduction* in cooling energy demand via liquid cooling. Additionally, liquid cooling allows high-density racks (over 50 kW/rack) that would be impractical with air cooling, enabling more compact and potentially more material-efficient data center designs.
2. **Immersion Cooling:** An extreme form of liquid cooling where servers are submerged in a dielectric fluid. This approach virtually eliminates the need for air conditioning units – heat is transferred to the fluid and then to a secondary water loop or coolant tower. Immersion systems can achieve PUE improvements by eliminating most fan power and enabling warm water cooling. Case studies report not only energy savings but also a reduction in hardware failure rates, since components operate in a stable, low-corrosion environment.
3. **Free Cooling:** Using outside ambient air or water to cool IT equipment when conditions allow. Many data centers now have economizer modes – when the outside air is cool enough, they circulate it directly (after filtering) to cool the servers, bypassing chillers. Alternatively, cooling towers or dry coolers reject heat to the outside without compressors. In cool climates, free cooling can be available > 50% of the time, drastically cutting electricity use. Even in moderate climates, designs that expand the allowable temperature/humidity range for servers (per ASHRAE’s guidelines) increase the hours of free cooling [20].
4. **Thermal Energy Storage:** Some facilities are adding thermal storage (like chilled water tanks or ice storage) to shift cooling load. The idea is to produce cold water or ice during off-peak times (or when renewable energy is plentiful), then use it to absorb heat during peak times, reducing chiller operation when grid power is dirtiest or most expensive. This strategy can also provide resiliency, as stored cooling can keep servers cool through short power interruptions.
5. **Advanced Airflow and Controls:** Beyond the hardware changes, advanced cooling also involves smarter control algorithms (often AI-based, as noted). Sensors throughout the data center feed into models that adjust cooling unit output, valve positions, and airflow in real time to eliminate hotspots with minimal energy. For example, machine learning control at Google’s data centers yielded a 40% reduction in cooling system energy by optimizing setpoints continually.

Each of these techniques contributes to lowering the **effective PUE**. Many new hyperscale data centers now report cooling PUE (sometimes called partial PUE for cooling) well under 1.2 thanks to combinations of the above methods. Importantly, reducing cooling energy not only cuts electricity use but also indirectly reduces carbon emissions and can allow more of a facility's power budget to go to IT equipment (doing useful work). From a carbon perspective, if a data center is powered by fossil-heavy electricity, cutting 1 MWh of cooling saves roughly 0.4–0.8 ton CO₂ (depending on grid intensity), which multiplied over thousands of hours yields substantial savings. Advanced cooling is thus a critical strategy in the path to net-zero data centers.

5.5. Modular and Edge Data Centers

The rise of edge computing—locating servers closer to end-users for low latency—is changing the landscape of digital infrastructure. Instead of solely massive centralized data centers, we now see many smaller “micro” data centers distributed geographically. While edge facilities help reduce network energy (shorter data transmission distances) and improve performance, they bring efficiency challenges. Smaller sites often cannot achieve the same PUE levels as large data centers, due to less efficient cooling (package HVAC units vs. industrial chillers), and lower economies of scale. However, **modular design** principles are being applied to make edge data centers as efficient as possible. Prefabricated modular data center units often come with integrated, optimized cooling and power systems tailored to the module's IT load, reducing over-provisioning.

These modules can be tested for efficiency in factories and use advanced cooling like in-row cooling or liquid cooling appropriate for their scale. Edge data centers also have the opportunity to utilize local renewable energy—for example, a rooftop solar array on a telecom exchange hosting an edge data center, or fuel cells providing clean power at a remote site. The key is to incorporate the same best practices in hardware efficiency, cooling, and energy sourcing that large facilities use, but on a smaller scale. Another aspect is **dynamic management**: since edge nodes might have variable load (peaking during local usage hours), operators can remotely power down or sleep servers during low-demand periods to save energy. Aggregating many edge sites, centralized AI systems can manage their workloads and energy use collectively, much like in a cloud region. Overall, while modular and edge data centers individually have a modest footprint, their proliferation means collectively they must not be ignored in efficiency efforts. Designing edge infrastructure with net-zero principles (efficient, renewable-powered, smart-managed) ensures that the expansion of the network's edge does not undercut the sustainability gains made at the core.

VI. POLICY AND CORPORATE ACTION TOWARD NET ZERO

Reducing the carbon footprint of digital infrastructure is not only a technical challenge but also a policy and corporate governance challenge. Governments, industry consortia, and companies themselves are all key players in driving the transition to net-zero emissions. This section discusses how regulatory incentives, corporate commitments, and collaborative initiatives are accelerating progress, and also examines the role of green building standards like LEED in promoting sustainable data centers.

6.1. Regulatory Incentives and Requirements

Policymakers can drive digital infrastructure decarbonization through a mix of incentives and regulations. A prominent tool is **carbon pricing** – imposing a cost on carbon emissions (via a carbon tax or cap-and-trade system) to encourage investments in cleaner energy. If data center operators must pay, for example, \$50 per ton of CO₂ emitted, they have a direct financial motive to reduce consumption or buy cleaner power. Recent analyses suggest that targeted carbon fees on data centers and crypto operations could meaningfully cut emissions; an IMF study estimated a \$0.032/kWh tax on data center electricity could align their growth with climate goals [11]. Beyond pricing, many jurisdictions are adopting **efficiency standards**.

These could mandate a maximum PUE or set an annual energy usage baseline that data centers must not exceed relative to their IT load. While such standards are still emerging, proposed rules in the EU contemplate PUE and temperature setpoint requirements for new data centers. Some cities and states have introduced **building energy codes** that include data centers, requiring efficient equipment and sometimes waste-heat recovery for large installations. Another regulatory approach is **transparency and reporting**. Governments (e.g., in the EU, as noted) are requiring data centers to publicly report energy and water usage, and GHG emissions. This transparency can drive competition to improve and allows regulators to benchmark performance. In regions like California, large facilities must already report under climate programs, and there are discussions to integrate data centers into broader emissions reporting schemes. On the incentive

side, governments offer grants or tax incentives for energy-efficient or green-certified data center projects. For instance, a state might provide a sales tax abatement on equipment if the data center meets **LEED certification** or uses a certain percentage of renewable energy. Such carrots can offset the cost of investing in more efficient technology. Finally, grid regulators are increasingly viewing data centers as part of the energy ecosystem – encouraging them to provide demand response or utilize on-site generation to support the grid (as the US DOE highlighted [7]).

In exchange, data centers may get favorable tariffs or infrastructure support. In summary, regulatory measures are tightening for data centers: the direction is toward greater efficiency, accountability, and integration with climate policy. Forward-looking operators are anticipating these changes by adopting higher standards voluntarily before they become mandatory.

6.2. Corporate Climate Commitments

The private sector, especially large tech companies, has taken a lead in setting ambitious climate and sustainability goals for digital infrastructure. Many of the world's largest data center operators have pledged carbon reduction or net-zero targets within specified timelines. For instance, Microsoft has committed to be **carbon negative by 2030**, [5] meaning it will remove more carbon than it emits, and to run on 100% renewable energy by 2025. Google aims to achieve **24/7 carbon-free energy** for its data centers by 2030, moving beyond annual offsets to hour-by-hour matching of renewable power [10]. Amazon's Climate Pledge targets net-zero carbon by 2040 across its operations (including AWS data centers), and in the interim Amazon is already the world's largest corporate renewable energy buyer [25] to reach 100% renewable by 2025. Meta (Facebook) has achieved net-zero emissions for operations and maintains 100% renewable energy sourcing as of the early 2020s [14].

These public commitments often exceed regulatory requirements and serve multiple purposes: they signal environmental responsibility to consumers and investors, and they set internal targets that drive innovation. The influence of corporate commitments is evident in industry progress. According to recent data, the portfolios of renewable energy contracts held by big tech companies rival those of utilities [25], as noted earlier, directly attributable to their self-imposed goals. Additionally, companies are investing in R&D for efficiency and novel technologies (like Microsoft's exploration of hydrogen fuel cell generators, Google's AI for cooling, etc.) to meet their targets. Another dimension is **supply chain engagement**: firms like Google and Microsoft are also pushing their suppliers (including hardware manufacturers) to reduce carbon footprint, effectively extending their climate commitments beyond their own data centers (Scope 3 emissions).

The cumulative effect is that corporate climate commitments by data center operators are **catalyzing innovation and setting de facto industry standards**. When a market leader demonstrates that net-zero operation is feasible, it creates competitive pressure for others to follow suit or risk being seen as laggards. Moreover, these commitments often involve joining international initiatives – for example, many tech companies joined the **Climate Neutral Data Centre Pact** in Europe or RE100 (a pledge for 100% renewable power). In summary, corporate commitments are a driving force pushing the envelope of what is achievable in sustainable digital infrastructure, accelerating timelines for decarbonization ahead of what regulation might require.

6.3. Collaboration and Knowledge Sharing

The challenge of greening data centers is complex, and collaboration across organizations magnifies the impact of solutions. Industry consortia and non-profits play a key role in disseminating best practices and developing standards. Organizations like **The Green Grid** (which introduced PUE and other metrics) and the **Uptime Institute** regularly publish survey results and guidelines that help operators benchmark themselves [9]. The Uptime Institute's annual Global Data Center Survey, for instance, provides insights into trends in energy reporting, outages, PUE, and sustainability initiatives, allowing the industry to track progress collectively. Another example is the **Open Compute Project (OCP)**, an initiative started by Facebook (Meta) that open-sources efficient data center and hardware designs. Through OCP, companies share designs for energy-efficient power supplies, server boards optimized for airflow, and even entire facility designs that anyone can use.

Collaboration is especially important for newer challenges like measuring **embodied carbon** (the CO₂ from construction and hardware manufacturing). Initiatives now are looking at standardizing how to report embodied carbon for data centers [14], so that operators can work with construction firms and equipment vendors to reduce it. In essence, collaboration accelerates the adoption of sustainable practices by ensuring that lessons learned in one project become shared knowledge

rather than proprietary advantage. As the sector seeks to hit ambitious climate targets in a short timeframe, this collective approach – “green together” – is proving to be a powerful tool.

6.4. Green Building Standards and LEED Certification

An increasingly important aspect of sustainable data centers is **green building design**, as recognized by certifications like LEED (Leadership in Energy and Environmental Design). LEED, developed by the U.S. Green Building Council, is a globally recognized rating system for environmentally friendly building design, construction, and operation. In recent years, LEED standards have been adapted specifically for data centers (LEED BD+C: Data Centers) to account for their unique features and loads [15]. Earning LEED certification (at Certified, Silver, Gold, or Platinum levels) signifies that a data center project has met stringent criteria across energy efficiency, water usage, site sustainability, materials, and indoor environmental quality.

For data centers, **LEED certification** drives a holistic approach to sustainability beyond just IT efficiency. Key LEED measures include: highly efficient HVAC systems, optimization of energy performance (often demonstrated via low PUE and use of economizers), integration of on-site renewable energy, use of sustainable building materials, water conservation (e.g., rainwater harvesting for cooling or using air-cooled chillers to minimize water use), and innovative design features like heat recovery. By pursuing LEED, data center builders must document and implement these features, which often leads to lower operating costs and reduced environmental impact.

Notably, several high-profile data centers have achieved **LEED Gold and Platinum** status, showcasing what is possible:

- **Apple's Maiden Data Center (North Carolina)** – One of the first enterprise data centers to earn **LEED Platinum**, it features an energy-efficient design and extensive on-site renewable energy. Apple's Maiden facility includes a large 20 MW solar farm and biogas fuel cells that together provide a significant portion of its power, enabling the site to be run on 100% renewable energy [23]. The building's design optimized lighting and cooling systems, and Apple reported a PUE near 1.2 for this facility, exceptional for its time. The LEED Platinum rating was achieved by excelling in categories like Sustainable Sites and Energy & Atmosphere (where it earned points for both on-site generation and high efficiency).
- **Citigroup Data Center (Frankfurt, Germany)** – This facility was the *first data center in the world to achieve LEED Platinum* certification [20]. Opened in 2012, it represented a “quantum leap” in energy efficiency at the time, with a PUE of < 1.2 achieved through innovative design [20]. The data center uses **70% free cooling** (outside air and adiabatic cooling) meaning chillers are bypassed most of the time, and it has a green roof and a prominent green wall irrigated by rainwater, reducing the urban heat effect and insulating the building [20]. Inside, efficient power distribution and partitioning minimized loss, and a novel water recycling system cut cooling tower water consumption by 30% [20]. These features not only earned LEED points but also translated to lower operating carbon emissions. Citigroup's achievement demonstrated that even financial sector data centers (traditionally very conservative about reliability) can implement cutting-edge sustainability without compromising performance.
- **Meta's Data Centers (various locations)** – Meta (formerly Facebook) has committed to LEED certification for all its data center buildings, achieving LEED Gold for 100% of its data centers by 2021 [14]. For example, the Gallatin Data Center in Tennessee is LEED Gold, supported by 100% renewable energy and designed to use **80% less water** than a typical data center of its size [21]. Meta's data centers focus on efficient cooling through outside air and evaporative techniques, and their custom server hardware (through OCP) is designed for optimal power efficiency. As a result, Meta reported industry-leading PUE figures (annual average of 1.09 in 2023 as noted) and very low water usage effectiveness (WUE of 0.20) [14]. These performance metrics, coupled with sustainable construction practices (91% of construction waste recycled, use of recycled materials in construction), fulfill and exceed LEED requirements, demonstrating that large hyperscale campuses can combine scale with sustainability [24].

VII. TOWARD NET-ZERO DIGITAL INFRASTRUCTURE: A ROADMAP

Efficiency by Design – Implement energy-efficient hardware and consider life cycle impacts during procurement. Plan system architectures to maximize utilization and avoid idle resources, using virtualization and consolidation to reduce the hardware footprint.

Data-Driven Operations – Employ AI for real-time monitoring, predictive maintenance, and intelligent workload distribution. Use advanced analytics to continuously optimize server utilization and dynamically adjust cooling for current conditions.

Sustainable Energy Sourcing – Procure renewable energy through PPAs or on-site generation (solar panels, wind, etc.), supplemented by energy storage solutions. Explore novel backup power options like battery systems or hydrogen fuel cells to replace diesel generators, further reducing carbon emissions during grid outages.

Transparent Reporting – Regularly measure and publish key sustainability metrics (PUE, CUE, WUE). Adopt and exceed existing standards (e.g., **LEED**, ISO 50001 energy management) to demonstrate accountability and progress. Public transparency builds trust and drives internal improvement.

Collaborative Innovation – Form partnerships across industry and with utilities and governments for R&D on efficiency and clean energy. Participate in consortia and initiatives (Open Compute Project, climate accords) that enable sharing of best practices and setting of higher standards collectively. Collaboration accelerates innovation and helps align efforts with broader grid decarbonization and climate goals.

By following this integrated roadmap, stakeholders can make concrete progress toward **net-zero digital infrastructure**, ensuring that the growth of the digital economy can be sustained within the limits of our climate objectives.

VIII. CONCLUSION

Digital infrastructure underpins the modern world, yet its energy footprint is significant and growing. This study has shown that with a holistic approach—encompassing efficiency improvements in hardware and cooling, a transition to renewable energy, intelligent management of workloads, supportive policies, and green building practices—data centers and cloud computing can drastically reduce their carbon emissions. The latest data reveals both cause for concern and optimism: while global data center energy use is rising due to new demands (like AI), efficiency innovations and renewable deployments are accelerating to counteract this growth. Concurrently, regulatory incentives and corporate net-zero commitments are catalyzing change, pushing operators to adopt best practices and report progress transparently. Achieving net-zero digital infrastructure is a formidable challenge that requires coordinated action across technology, industry collaboration, and policy enforcement.

However, the examples highlighted—from ultra-efficient, LEED-certified data centers to AI-optimized operations—illustrate that the vision of sustainable computing is attainable. As our reliance on cloud services and digital connectivity deepens, it is imperative that we continue on this path. Through persistent innovation and shared commitment, the industry can harmonize the digital future with global climate objectives, ensuring that the next era of computing is not only faster and smarter but also cleaner and more responsible.

REFERENCES

- [1] International Energy Agency (IEA). (2020). Data Centres and Data Transmission Networks. Paris: IEA.
- [2] Koomey, J. (2011). Growth in Data Center Electricity Use 2005–2010. Analytics Press, 1–19.
- [3] Shehabi, A., et al. (2016). United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory.
- [4] Beloglazov, A., & Buyya, R. (2012). Optimal algorithms for energy-efficient dynamic consolidation of VMs in cloud data centers. *Concurrency and Computation*, 24(13), 1397–1420.
- [5] Microsoft. (2020). Microsoft will be carbon negative by 2030. Microsoft News Center.
- [6] Patel, C.D. (2003). A Vision of Energy Aware Computing from Chips to Data Centers. HP Labs.
- [7] U.S. Department of Energy (DOE). (2024). 2024 Report on U.S. Data Center Energy Use. Lawrence Berkeley National Lab.
- [8] Gooding, M. (2024). Global data center electricity use to double by 2026 – IEA report <https://www.datacenterdynamics.com/en/news/global-data-center-electricity-use-to-double-by-2026-report/>
- [9] Uptime Institute. (2022). Global Data Center Survey Results. (Press Release) <https://uptimeinstitute.com/about-ui/press-releases/2022-global-data-center-survey-reveals-strong-industry-growth>
- [10] Google. (2020). Carbon-Aware Load Shifting at Google Data Centers. Google Cloud Blog <https://www.datacenterknowledge.com/hyperscalers/google-says-some-computing-its-data-centers-do-can-wait-for-sunshine>
- [11] Hebous, S., & Vernon-Lin, N. (2024). Crypto and AI – Charting their growing carbon footprint. IMF Blog <https://www.imf.org/en/Blogs/Articles/2024/08/15/carbon-emissions-from-ai-and-crypto-are-surging-and-tax-policy-can-help>
- [12] Chandrakant, P., & Wang, Y. (2021). Immersion cooling of data centers: A case study. *IEEE Trans. on Components, Packaging, and Manufacturing Tech.*, 11(2), 197–206.
- [13] Climate Neutral Data Centre Pact. (2020). Pact Memorandum of Understanding. Brussels, EU.

- [14] Meta Platforms. (2023). Meta Sustainability: Data Centers <https://sustainability.atmeta.com/data-centers/> . Meta Sustainability Website.
- [15] U.S. Green Building Council (USGBC). (2023). Applying LEED to Data Center Projects. USGBC Knowledge Base <https://support.usgbc.org/hc/en-us/articles/12154267763987-Applying-LEED-to-data-center-projects>
- [16] Data Centres and Data Transmission Networks <https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks>
- [17] How much energy do data centers use? <https://davidmytton.blog/how-much-energy-do-data-centers-use/>
- [18]] DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers <https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers>
- [19] Global PUEs — are they going anywhere? <https://journal.uptimeinstitute.com/global-pues-are-they-going-anywhere/>
- [20] Sustainable and resilient data center design: the world's first LEED Platinum certified data center <https://www.arup.com/projects/citigroup-data-center-frankfurt-am-main/>
- [21] Henrico Data Center's Post <https://www.facebook.com/HenricoDataCenter/posts/we-are-proud-to-share-that-buildings-5-and-6-on-our-henrico-data-center-campus-a/1267278047922972/>
- [22] Amazon's new cloud computing chip, AWS Graviton4, is now generally available. Take a look back at its evolution. <https://www.aboutamazon.com/news/aws/graviton4-aws-cloud-computing-chip>
- [23] How Many Data Centers Are LEED Platinum? update from February 2012 <https://www.itprotoday.com/macros/how-many-data-centers-are-leed-platinum-70431ba1-dff1-48e7-9b59-f456e63c1e79>
- [24] Meta: Supporting energy needs, reducing life cycle impacts <https://www.usgbc.org/articles/meta-supporting-energy-needs-reducing-life-cycle-impacts>
- [25] Cloud Titans Were the Largest Buyers of Renewable Energy in 2021 <https://www.datacenterfrontier.com/featured/article/11427604/cloud-titans-were-the-largest-buyers-of-renewable-energy-in-2021>
- [26] Datacenter companies continue renewable buying spree, surpassing 40 GW in US <https://www.spglobal.com/market-intelligence/en/news-insights/research/datacenter-companies-continue-renewable-buying-sprees-surpassing-40-gw-in-us>
- [27] Google cuts on-site carbon emissions but struggles to contain data center electricity use <https://www.facilitiesdive.com/news/google-environmental-report-emissions-increases-data-centers/721112/>



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details