

Verifiable Sustainability in Data Centers

Syed Rafiul Hussain  | Pennsylvania State University
Patrick McDaniel  | University of Wisconsin–Madison
Anshul Gandhi  | Stony Brook University
Kanad Ghose and Kartik Gopalan  | Binghamton University
Dongyoon Lee  | Stony Brook University
Yu David Liu | Binghamton University
Zhenhua Liu, Shuai Mu, and Erez Zadok  | Stony Brook University

The current techniques and tools for collecting, aggregating, and reporting verifiable sustainability data are vulnerable to cyberattacks and misuse, requiring new security and privacy-preserving solutions. This article outlines security challenges and research directions for addressing these requirements.

Sustainability is the practice of performing human activities in ways that do not leave lasting harmful effects. Unfortunately, harm to the planet is clearly growing,^{16, 17} whether the effects are direct (e.g., emissions caused by transportation, farming, or manufacturing) or indirect (e.g., carbon emissions due to electricity consumed by data centers and even the energy and materials used for manufacturing servers and other devices). Humans as a species have understood that sustainability is important to both future generations and the global quality of life. Yet, we have had only sporadic and uneven adoption of sustainable practices, and up to 98% of sustainability initiatives fail to meet their goals.¹⁸ The impacts of a lack of sustainability have led to, among many other factors, climate change, widespread pollution of the oceans, sea bottom desertification, acidification of land and water, ozone loss, desertification, and loss of biodiversity. Failure to address this lack of sustainability now will create long-term problems for future generations.

Today, achieving the goals of sustainability requires the honest best efforts of humans and apparatus to measure aspects of the system under regulation. Yet, those efforts often fail when bad actors bypass or cheat sustainability systems. For example, the car company Volkswagen installed emissions software on roughly

11 million cars worldwide that misled the U.S. Environmental Protection Agency (EPA) about emissions when under test.¹⁹ Volkswagen was eventually caught, fined billions of dollars, and required to recall vehicles and pay financial settlements—but only after the vehicles had polluted for nearly a decade.

One area with unprecedented impact on our world is the use of computation and, in particular, data centers. With the alarming rise of computation and the pervasive use of artificial intelligence (e.g., ChatGPT),²⁰ data centers have many negative impacts on the environment, caused by energy use, hardware manufacturing and disposal, building maintenance, water usage, and other factors. Indeed, a recent study showed that over 2%–4% of all energy used worldwide was by data centers.⁹ The current practice of reporting sustainability information in data centers is, however, mired with “greenwashing,” where the true carbon footprint of a data center is artificially reduced via the purchase of energy or certificates from green generation sources²¹ or by paying other entities to be sustainable. This signifies a lack of transparency and accountability that hinders efforts to address and mitigate the environmental consequences associated with data centers. Such issues are pervasive, as they extend beyond data centers and permeate various industries, including food, manufacturing, and telecommunication systems.

Digital Object Identifier 10.1109/MSEC.2024.3372488

The lack of accountability and transparency to address sustainability is primarily rooted in the absence of complete and verifiable sustainability data and metrics.²² Comprehensive and fine-grained sustainability metrics are critical to identify performance bottlenecks (e.g., the impact of an application's code or library on sustainability), diagnose security issues, detect anomalous sustainability activities, provide a reliable audit trail of carbon consumption, ensure accurate and precise accountability and compliance benefits (e.g., accurately identify entities that make changes or perform certain actions), and optimize system performance.²³ Therefore, a necessary first step for any sustainable computing approach is the ability to measure comprehensive sustainability metrics or cost functions from all possible sources of carbon consumption and energy spent in the entire life-cycle of the computing equipment: production, delivery, and disposal; these are referred to as “embodied energy.” However, it has been found that it is difficult to determine accurate sustainability metrics because the sources are too many, untrustworthy, disconnected, or incompatible. Further, there is no way to combine the data in a meaningful way that will not compromise the privacy of users or service providers.²⁴ For example, there are dozens of different ways to calculate data on global data center energy consumption that are based on public and private data, each resulting in an assessment that is often contradictory with others.²⁵ Hence, we have at best a vague idea of the impact that, for example, data centers have on our environment. Even when attempts are made to collect and combine sustainability metrics from disparate sources, privacy concerns and exposure of sensitive user data or service providers' proprietary algorithms are often ignored, resulting in poor incentives for users or service providers to opt for accountable sustainability systems. Researchers and organizations trying to understand and create sustainable systems often refer to the *sustainability data gap*. The inability to collect and verify accurate, complete, and timely data on the environment in a privacy-preserving fashion is slowing, and in some cases prohibiting, the adoption of sustainable systems and practices. To make matters worse, market forces and human greed, as we observed earlier, often work against the goals of sustainability.

In the context of data centers, which is the primary focus of this article, the infrastructures used to measure and maintain operational sustainability (i.e., environmental footprints within a data center) are inherently adversarial: because users of technology (e.g., data center users) have an incentive to cheat, the apparatus must strive to ensure that systems continue to function correctly in the face of actors attempting to thwart the collection of sensitive sustainability footprints and the enforcement of corresponding security and privacy

policies. Hence, it is imperative that the environmental footprint of data center operations can be verified by interested third parties (e.g., the EPA,²⁶ citizen scientists, and the public).

This article, therefore, identifies the security issues in the sustainability data pipeline consisting of data collection, storage, aggregation (or other processing), reporting, and use in situ. More specifically, we examine threat landscapes and a wide range of security challenges to build verifiable sustainability within data centers, highlighting the urgent need to address these threats. Furthermore, we explore a variety of promising research directions that will yield novel and practical solutions to combat these security challenges in sustainable data centers and mitigate the risks associated with such threat landscapes. Some of our proposed security challenges and solutions also apply to other industry segments: manufacturing, airlines and transportation, industrial-scale farming, and more.

Sustainable Systems and Focus on Data Centers

There are several systems (or industries) whose unsustainable operations pose a grave threat to the environment. For example, sustainability concerns are important across a wide industry segment, such as livestock farming, automobiles, airlines, manufacturing, energy generation, and transportation as well as infrastructure construction and management (e.g., those applicable to buildings and roadways). Data centers are particularly significant due to their substantial energy consumption and environmental impact. Moreover, data centers play a vital role in supporting many industries and services that rely on digital infrastructure, making their sustainability practices even more critical.

Operations within data centers already contribute significantly to the global carbon footprint.¹² The rise in popularity of resource-intensive big data, artificial intelligence, cryptocurrency, and machine learning workloads is poised to make data center operations even more unsustainable.²⁷ Estimates suggest that data centers are already responsible for about 2%–4% of the total greenhouse emissions; that is equivalent to the emissions of the entire airline industry.²⁸ Worse, this figure for data centers is soon expected to increase to 5%–7% with the emergence of large language models (LLMs), such as GPT-4, and applications based on LLMs, imposing a heavier toll on the environment.²⁹

This article, therefore, specifically focuses on sustainability in data centers. Even though data centers' energy consumption can be significant, their efficiency through shared resources—the ability to integrate renewable energy and optimize computing power—can potentially reduce their overall footprint

compared to distributed on-premise solutions, such as edge servers or private clouds. Also, the sheer volume of computing resources to process a wide range of data, optimized cooling systems, and advanced energy consumption equipment in data centers enables us to obtain a comprehensive view of their environmental impact as opposed to on-site alternatives. This could be achieved by a fine-grained approach to measuring energy consumption/carbon, which requires security primitives, as it could become a criterion for optimization, fine-grained diagnosis, and decision making, similar to financial cost, in the long term. This effort is also necessary to complete net zero and aligns with the aspirations of the largest IT companies.³⁰ Moreover, such transparent monitoring and verifiable audits of energy consumption will help data centers and service providers establish trust with customers, investors, and regulators. Therefore, in this article, we specifically focus on sustainability in data centers.

Existing practices in data centers on reporting or advertising sustainability data are often fraught with greenwashing; as a result, the true carbon usage of a data center is hidden. Similar greenwashing practices have also been observed in other sectors, including autonomous vehicles³¹ and telecommunication industries.³¹ Such deceptive approaches undermine the transparency and credibility of sustainability claims, making it difficult for stakeholders to make informed decisions. The European Union's (EU's) Corporate Sustainability Reporting Directive mandates that by 2024, corporations have to report nonfinancial sustainability information precisely and clearly; this will also apply to data center operators within the EU. There is some consensus among data center operators on reporting data center sustainability information and metrics accurately, at least within the EU and the Asia-Pacific region.³² In the United States, we also see the beginnings of directives similar³³ to the EU's, but details are still emerging.

Coarse-grained accounting may be possible using recent carbon measurement prototypes, such as Power API (<https://powerapi.org/>), Kepler (<https://sustainable-computing.io/>), and Scaphandre (<https://github.com/hubblo-org/scaphandre>); however, they are still in their infancy and are not designed to provide any support for the verifiability of the generated carbon footprints by regulatory agencies or to ensure the privacy of users' sustainability data. The lack of such security and privacy guarantees, as outlined in the "Why Is Sustainability a Security Problem?" section, can be exploited by malicious entities (e.g., data center and service providers) to bypass carbon compliance, evade taxes, inflict financial losses on rival companies, cause over/under-billing of customers, steal sensitive user data and proprietary models, and contribute to environmental hazards.

Hence, security measures are indispensable for independent audits to provide an objective and verifiable assessment of data centers' sustainability claims, prevent greenwashing, and ensure accurate reporting of carbon emissions. Also, this transparency and accountability build trust with stakeholders like customers, investors, and the general public, who are increasingly concerned about the environmental impact of data centers.

Data Center Architectures

Data centers can be of different types based on size, purpose, and the services they offer. For instance, enterprise data centers are operated by individual organizations to manage and store their own data and IT infrastructure, whereas edge data centers are smaller facilities located closer to end users to reduce latency and enhance the performance of edge computing applications. This article, however, primarily focuses on colocated, hyperscale, distributed, or cloud data centers. Colocated data centers (e.g., Equinix and Digital Realty) provide space, power, and cooling for servers owned by different organizations, or tenants, promoting the sharing of facility resources and physical infrastructures. Hyperscale and cloud data centers, such as Google, Amazon, and Microsoft, deliver services like infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) over the Internet and can handle massive amounts of data and traffic. IaaS providers abstract and virtualize the underlying physical IT infrastructure and create isolated virtual environments, thus enabling end users and customers to run applications, store data, and utilize physical resources provided by data center providers. PaaS providers offer platforms (e.g., development tools, middleware, and database and deployment services) as a service to SaaS providers to streamline application development, deployment, and management, facilitating faster and more efficient processes. SaaS providers grant users instant access to software applications, data storage, access control, application programming interfaces (APIs), and integration, eliminating the need for businesses to invest in and maintain hardware and software, thereby reducing overall IT infrastructure costs.

Why Is Sustainability a Security Problem?

Ensuring the accuracy and credibility of sustainability metrics, as well as empowering audits by regulatory agencies, require guaranteeing the trustworthiness and comprehensiveness of not only the carbon footprints of data center equipment but also the embodied energy throughout the entire lifecycle of computing equipment. Although some external information, such as that for renewable energy, energy credits, or supplied water,

can be authenticated via trusted third parties, sustainability metrics in data centers require the authenticity, confidentiality, integrity, and availability of data collected, processed, stored, and used locally within a data center.⁴ However, unlike traditional cloud computing systems, where the focus is primarily on security and privacy of user applications and data, collecting and measuring data center activities that impact humans and the environment in a verifiable and privacy-preserving manner presents a diverse set of new security challenges. Most of these challenges are primarily based on sustainability data, reliability of equipment, and cleanliness of energy sources across both the digital and physical worlds. Unfortunately, no prior research has investigated the threat landscape of sustainable data centers nor attempted to provide any techniques or tools that directly allow authentication of operational sustainability metrics within a data center to preserve the privacy of users' or operators' sustainability data.

Also, we note that a key distinction with sustainable data, as opposed to regular data in cloud infrastructure, is that they are generated independently of user intent, resulting in reduced trust guarantees. The critical factor here is to prevent users from misrepresenting their emissions. Therefore, these data must be generated, collected, and aggregated in a manner that is tamper resistant, akin to a physical value, ensuring that they are nearly impossible for anyone to manipulate and do not expose any sensitive information about users. In a nutshell, this threat model is different from most common data, as the trust has to be minimal.

It is thus imperative to ensure the security of 1) data collection processes, 2) the process of generating verifiable and easily auditable sustainability metrics, and 3) the storage of all pertinent information. Hence, while being indispensable for protecting the environment and our planet, we have found and argue that the current sustainability practices—through self-reporting, best-effort measurement, and anything less than complete verifiable control of sustainability—will fail.

Threat Models in Data Centers

The trust assumptions and threat models for sustainable data centers may vary widely based on data center type (e.g., multitenant and hyperscale versus enterprise data centers), service models offered by the data center or the tenants, and any other specific requirements. In general, the threat models for a colocated, hyperscale, or cloud data center's sustainability can be primarily derived with respect to four entities: 1) data center providers, 2) tenants or service providers, 3) users, and 4) third-party observers (e.g., regulatory agencies), leading to the following adversarial capabilities:

- \mathcal{A}_1 : Here, data centers provide misleading or false sustainability data to attract end users or third-party service providers.
- \mathcal{A}_2 : Data centers or tenants providing IaaS, PaaS, and SaaS, often characterized as honest but curious, may attempt to learn proprietary or sensitive data of their users and exfiltrate it to third parties.
- \mathcal{A}_3 : Data center or service providers have access to their users or tenant's sustainability data and can be inherently malicious, exploiting this information to harm users or learn proprietary information that would benefit competitors or harm their tenants/consumers.
- \mathcal{A}_4 : Tenants (i.e., service providers), on the other hand, can also subvert the security and privacy of other colocated tenants' resources and the facilities provided to them.
- \mathcal{A}_5 : To make matters worse, resources (hardware and software) served by tenants (e.g., IaaS, PaaS, or SaaS) within a data center can also be compromised and controlled by external attackers who are nation-states or rival organizations offering similar services. This is possible due to system/service misconfigurations, insecure communication protocols, inadequate access controls and isolation of shared and physical resources, and vulnerabilities in the hardware, software, or other components of the service providers' supply chains. For example, benign and unsuspecting data center providers often use virtual machines (VMs) or containers created by IaaS providers that are loaded with backdoors or malware illegitimately reading/writing sensitive carbon footprint data.
- \mathcal{A}_6 : The other key entities in data centers (i.e., users or customers of a tenant) can also be considered malicious. This is because a user's job (e.g., a process) running in a data center may attempt to gain unauthorized access to read or modify other jobs' code and data and thus affect the sustainability data produced by other jobs. For example, a malicious process of an end user may add unaccounted read/write operations² to users' jobs, which can inflate users' carbon footprints, leading to overbilling the victim customers. Such carbon footprint inflation can also be achieved by violating the integrity of the sustainability metrics (e.g., code or data)² or by manipulating the system traces and logs—the evidence trail of carbon consumption by the compromised VMs or malicious processes in data centers. Similarly, compromised data center providers may exploit the same and use similar malicious processes to report false carbon footprints to regulators to evade high carbon taxes or regulations.² Users may also try to launch attacks [e.g., denial of service (DoS)] against other users or the tenant who owns that service as well as another tenant or its users in

the same data center. Users may also strive to obtain higher levels of service than they are allocated and thus mislead service providers about their carbon footprint. Various surfaces can be utilized by users to attack the tenant, including hypervisors, VMs, APIs, and web services.

Last but not least, third-party observers (e.g., regulatory agencies) are tasked with verifying the footprint reported by the data center and service providers in the process of executing policy or oversight (e.g., by comparing sustainability costs reported by cloud operators, users, and utilities).

- \mathcal{A}_7 : But even these observers may be honest but curious, government or law enforcement agencies performing surveillance, or untrusted, as they could collude with others to mislead reporting, may have rogue insider elements within a data center, and may even be under political or other pressure to “fudge” or misrepresent the data.

Differences with other systems. Although there are some similarities between data centers and Internet of Things (IoT) and enterprise IT systems regarding data collection, verifiability, and storage, the key distinction lies in the threat model between these systems. For instance, in most IoT settings, users, being the owners of their homes and devices, do not tamper with devices to generate false data. The users also generally trust trigger-action platforms capable of storing sensor data, as those platforms are key enablers of automation. In most cases, third-party smart apps (i.e., trigger-action rules) or external attackers are untrusted, as they are the primary attack vectors. Another notable distinction with sustainability data in data centers, as compared to regular IoT data, is that the sustainability footprints recorded by physical and virtual infrastructures (e.g., power generators, cooling systems, VMs, and hypervisors) are shared across mutually untrusted stakeholders. This gives rise to privacy concerns, which inherently differ from those in IoT or enterprise IT systems, where multiple users sharing the same physical environment (e.g., smart home and building) are mutually trusted. Hence, one user’s IoT activities are not considered sensitive/private to another user in the same home/building.

Security Challenges for Sustainability

Due to the complex design of data centers and intricate interactions among their stakeholders, it is necessary to characterize and address diverse security threats to the sustainability pursuit of data centers. Next, we discuss some critical security challenges for a data center aiming for sustainability, as summarized in Table 1. Note that the nature of threats will be different for different

sustainable systems (e.g., transportation and manufacturing) based on trust assumptions.

Evasive carbon offset techniques (C1). Data centers and large corporations often trade a known amount of carbon emissions with an uncertain amount of emission reductions to claim carbon neutrality (e.g., by investing in forestation elsewhere). This practice, also called *carbon crediting* or *climate crediting*, has been in place for decades. It is often exploited by large corporations, as it is extremely difficult, if not impossible, to track and verify whether the amount of emissions balances out the amount of reductions. Often, renewable energy credits (RECs) are used to offset the carbon footprint of a data center via the purchase of energy credits from a green energy generator. Similarly, power purchase agreements (PPAs)³⁴ are used to have a data center operator finance the installation of a green energy-producing farm (owned and managed by an independent party) to provide green energy to the data center over a long-term period covered under the PPA. For both RECs and PPAs, the authenticity of green energy is, however, often kept out of sight of users. Therefore, the lack of authentication, accountability, and transparency enables corporations (\mathcal{A}_1) to make false claims about the energy source while appearing in public to support sustainability efforts.

Lack of integrity of carbon emission sources (C2). According to threat model \mathcal{A}_3 , sensors and devices [e.g., power distribution units (PDUs)] used for tracking sustainability data can be tampered with by their owners, i.e., untrusted data centers, IaaS providers, or physically colocated \mathcal{A}_4 tenants, to either misreport to regulatory agencies or overcharge customers. Such false reporting by \mathcal{A}_4 tenants can cause a data center operator to readjust resource allocations/scheduling unnecessarily to adversely affect the data center’s sustainability footprint. In a similar vein, these sensors and devices can also become compromised by external attackers (\mathcal{A}_5), due to unintentional vulnerabilities or intended backdoors in their hardware, firmware, and software.³⁵ As a result, by taking control of those sensors and devices, attackers may reduce authenticity and forge carbon footprints to cause over/underbilling of customers by forging/manipulating carbon consumption records. Attackers may also generate false sustainability data or manipulate cooling systems to disrupt sustainability operations.³⁶ Similar kinds of sustainability data forgery attacks can also be carried out if there are vulnerabilities in the communication protocols (e.g., lack of authentication and replay protection) between sensors and the sustainability data aggregators gleaned carbon footprints from multiple such sensors. Due to such malicious actions,

additional water and electricity would be required to cool the targeted data center, resulting in an increased carbon footprint, higher operational costs, and disruption of sustainability efforts.

Inadequate access control and information flow control (C3). While resource sharing in data centers offers cost efficiency, it requires robust isolation techniques to prevent unauthorized access to tenants’ sensitive data. The lack of fine-grained and dynamic access control (such as discretionary access control, mandatory access control, or combinations thereof), adequate resource isolation, and information flow control

measures may allow attackers (\mathcal{A}_3 and \mathcal{A}_4) to obtain unauthorized access to sensitive sustainability data, potentially leading to data breaches, privacy violations, and other security issues. Furthermore, sustainability data can also be illegitimately tampered with by malicious user processes (\mathcal{A}_6) or compromised system processes (\mathcal{A}_5). Malicious processes may obtain unauthorized (read/write) access to sensitive resources (e.g., databases or protected memory regions storing sustainability data and states) by exploiting vulnerabilities in the access control policies.³⁷ As a result, regular sustainability operations are likely to be disrupted, which may cause the system to produce unwarranted

Table 1. Threats and security challenges for the sustainability of data centers and potential research directions.

ID	Threat Model	Vulnerabilities and Security Challenges	Impacts	Possible Ideas for Solutions
C1	Untrusted: data center operators (\mathcal{A}_1) Trusted: other stakeholders	Evasive carbon offset techniques and lack of authenticity, accountability, and transparency allow data center providers (\mathcal{A}_1) to trade a known amount of carbon emissions with an uncertain amount of carbon reductions.	Tax evasion, financial loss, and environmental hazards	Verifiable footprint collection (the “Verifiable Footprint Collection Architecture” section)
C2	Untrusted: data center operators (\mathcal{A}_3), tenants (\mathcal{A}_4), external attackers (\mathcal{A}_5) Trusted: other stakeholders	The lack of integrity (tamperproof guarantee) of carbon emission sources allows malicious providers (\mathcal{A}_3), physically colocated tenants (\mathcal{A}_4), or external attackers (\mathcal{A}_5) to forge, tamper with, or misreport carbon usage. ¹⁴	Cause over/underbilling of customers by tampering with carbon usage, evade regulatory agencies by misreporting low carbon emissions	Verifiable footprint collection (the “Verifiable Footprint Collection Architecture” section)
C3	Untrusted: data center operators (\mathcal{A}_3), tenants (\mathcal{A}_4), external attackers (\mathcal{A}_5), users (\mathcal{A}_6) Trusted: other stakeholders	Inadequate access control or information flow control mechanisms may allow attackers (\mathcal{A}_3 – \mathcal{A}_6) to access and tamper with databases storing carbon footprint trails.	Exposure of users’ private data, such as location, behavior, and intellectual properties	Verifiable footprint collection (the “Verifiable Footprint Collection Architecture” section)
C4 and C6	Untrusted: data center operators (\mathcal{A}_2), tenants (\mathcal{A}_4), users (\mathcal{A}_6) Trusted: other stakeholders	Disclosure of sustainability metrics to honest but curious (\mathcal{A}_2) or malicious service providers (\mathcal{A}_4) and users (\mathcal{A}_6), due to inadequate access control, cryptographic protections, or side channel vulnerabilities	Exposure of users’ private data, such as location, behavior, and intellectual properties	Privacy-preserving footprint collection and aggregation (the “Privacy-Preserving Footprint Collection,” “Privacy-Preserving Footprint Aggregation,” and “Public Sustainability Ledgers” sections)
C5	Untrusted: data center operators (\mathcal{A}_3), tenants (\mathcal{A}_4), external attackers (\mathcal{A}_5), users (\mathcal{A}_6) Trusted: other stakeholders	Cryptographic flaws and software vulnerabilities may allow attackers (\mathcal{A}_5) to forge proof of carbon usage.	Financial loss and disruption of data center operations	Verifiable footprint collection (the “Verifiable Footprint Collection Architecture” section)
C7	Untrusted: data center operators (\mathcal{A}_1), tenants (\mathcal{A}_4), regulators (\mathcal{A}_7), users (\mathcal{A}_6) Trusted: government	Multiple parties, such as providers (\mathcal{A}_3) and users (\mathcal{A}_6) or providers (\mathcal{A}_3) and regulators (\mathcal{A}_7), may collude to misreport carbon usage.	Tax evasion, financial loss, and environmental hazards	Verifiable footprint collection (the “Verifiable Footprint Collection Architecture” section)

carbon footprints, including a neutral footprint. Tampering with sustainability data by attackers (e.g., malicious service providers or malicious users) may result in overcharging legitimate users of the system (such as a data center), undercharging malicious users attempting to evade sustainability costs, or damaging the reputation of competing service providers. Attackers may also induce carbon exhaustion attacks on other users by misreporting carbon consumption or evading compliance checking of regulatory agencies by misreporting low carbon emissions when operating in test mode (similar to Volkswagen's scandal³⁸).

Sensitive information disclosure (C4). Collecting fine-grained sustainability data from disparate carbon sources (e.g., sensors and PDUs) to monitor and diagnose sustainability activities may also disclose sustainability metrics to service providers (\mathcal{A}_2) and other users. Such unauthorized footprint exposure will violate the privacy of users' data, location, behavior, and intellectual properties, such as proprietary scheduling techniques, which are factors used for competitive pricing for different service categories.^{7, 10} Unauthorized access to footprint data can enable an adversary to initiate DoS attacks (\mathcal{A}_4 and \mathcal{A}_6) on cotenants and thus prevent cotenants from realizing a desired sustainability target.

Cryptographic flaws and software bugs (C5). The ability of a sustainable system to provide carbon footprint proof to users and regulators is essential for ensuring the trustworthiness of the system. Such proof of footprint should be built with cryptographic constructs. However, flaws in the integration of cryptographic constructs with complex data center systems (e.g., using weak cipher suites^{37, 39}) or flaws in the software⁴⁰ (\mathcal{A}_5) may fail to generate unforgeable and accurate proof of consumption, enabling an attacker to drop, modify, replay, and inject fake footprints of carbon. This can disrupt the operations of sustainable systems.

Side channels in sustainability (C6). Due to shared hardware resources, colocated tenants' servers, and poor isolation among different processes running on the same hardware in data centers, side channel vulnerabilities¹³ (e.g., page faults, cache misses, power, and timing channels) may allow a malicious process (\mathcal{A}_4 and \mathcal{A}_6) to observe or tamper with carbon footprint patterns of other users' jobs/applications running on the same hardware. Such side channels allow an attacker to not only fingerprint the data traffic of other users but also extract the cryptographic keys or other confidential information of a user application by looking at the use of sustainability metrics.¹³ Attackers (\mathcal{A}_3 and \mathcal{A}_5) can exploit such sensitive information to blackmail or

embarrass other users/competitors (e.g., to force a competitor's stock to drop or short sell such stock).

Collusion for evasion (C7). Infrastructure providers (\mathcal{A}_7) and PDU providers (\mathcal{A}_3) may collude to misreport carbon footprints to regulators and users and thus may evade regulatory agencies. Such collusion attacks can be of different combinations, as infrastructure providers depend on third-party software and hardware vendors, which may also collude with each other for malicious purposes.

Research Directions for Securing Sustainable Data Centers

Although many solutions¹⁵ have been designed for data center security, most of them are not directly applicable to counter the security and privacy challenges toward sustainability, as discussed in the "Why Is Sustainability a Security Problem?" section. Therefore, we must develop technologies that will help build secure and trustworthy sustainable systems. Particularly, we must develop primitives that allow domain experts to construct and operate sustainable systems and verify the results. Next, we lay out several potential research directions for improving sustainability in data centers through security.

Verifiable Footprint Collection Architecture

One of the most important elements of a sustainable system is its ability to promote the responsible use of system resources, such as complying with carbon emission restrictions/taxes. However, claims of carbon usage must be accompanied by infrastructure that demonstrates a verifiable footprint to the public and regulatory organizations. This calls for architectures and systems that can collect publicly readable and verifiable sensor readings in adversarial settings. It is essential that these systems have the ability to scale seamlessly from small low-energy devices to larger enterprise-level data centers. The system architecture should have the ability to generate tamper-resistant proofs of carbon consumption that are unforgeable, accurate, and securely retrievable by authorized parties (which might include the public) in adversarial deployments. Furthermore, to provide higher security assurance, the design and implementation of these systems must be formally verified.

Potential solutions. Developing such a framework poses key challenges, including the need to establish and preserve a root of trust by using trusted hardware, such as the Trusted Platform Module, to secure a data center's carbon footprint measurement components. A trusted path should be established from the secure hardware up to the module that collects all the relevant metrics of a job and

further up to the component that verifies the accuracy of the reported metrics. This trusted path will be capable of producing tamperproof evidence of sustainability cost metrics using cryptographic proof systems.

One potential solution to ensure the security of sustainability-related components is to use a hardware-based trusted execution environment (TEE), such as ARM TrustZone, Intel SGX, AMD SEV, and Keystone. TEEs are deployed in nearly every commercial processor sold today and are the de facto standard to provide a tamper-proof execution environment that preserves the integrity and confidentiality of data and execution.² These environments provide isolation guarantees needed to certify that metric data are collected and reported accurately, even in the presence of malicious applications, operating systems (OSs), or hypervisors. A sustainability collector (see Figure 1) running in a TEE will securely collect the utilization details of a bare-metal, virtualized, or containerized job. The gathered metrics will create a comprehensive timeline of user-, system-, and process-oriented carbon footprints, culminating in a sustainability provenance record for the cloud. The sustainability collector will securely report the metrics to a sustainability certification agent, which will produce lightweight cryptographic proofs that empower third-party regulators and users to independently verify the claimed consumption.

Note that any flaws in the design or implementation of sustainability-related components, e.g., measurement or collection code running within TEEs and owned by respective TEE hosting entities (i.e., data center operators or service providers), may introduce new security challenges. For instance, attackers may exploit such flaws and bypass the tamperproof guarantees of the code. Therefore, it is crucial to ensure high-security assurance of these components through formal analysis before they are deployed. Also, the physical machines or VMs hosting the measurement code within TEEs and regulatory agencies need to verify during runtime the integrity of the trusted path from the secure hardware to corresponding TEEs periodically or when there are major changes (e.g., write operations) in the system.

Another potential concern is that current TEE platforms might lack adequate privileges to monitor the carbon or resource consumption of workloads that execute outside of a TEE. This might necessitate new hardware support for TEEs to allow secure monitoring of external workloads, including the host OS or hypervisor.

One possible alternative to TEEs is to explore the use of add-on monitoring hardware, akin to smart network interface cards (SmartNICs), that can collect sustainability metrics from outside the host. For example, AWS Nitro (<https://aws.amazon.com/ec2/nitro/>) enables SmartNICs to monitor and manage VM allocation

and scheduling while being technically “outside” the host OS. Similarly, sustainability-related components could potentially run on such add-on custom hardware with the necessary privileges to gather data from the host without being vulnerable to compromise by the host. Finally, sustainability data must be isolated from other workloads running on the same machine, providing protection against unauthorized access and tampering.

Privacy-Preserving Footprint Collection

Fine-grained sustainability data collected through disparate carbon sources, such as sensors and PDUs, in an unregulated manner may induce unintended disclosure of sensitive data. The exposure of sustainability records would otherwise break users’ privacy, data, location, behavior, and intellectual properties, such as proprietary scheduling techniques, trained machine learning models, and factors used for competitive pricing for service classes.^{7, 10} Also, attackers may attempt to tamper with sensor data before they are aggregated, which can lead to incorrect or misleading results. This can be especially problematic in safety-critical applications, such as autonomous vehicles or medical devices.

Potential solutions. In concert with the verifiable sustainability data collection architecture, differential privacy (DP) or local DP can be used as a probabilistic solution for privacy-preserving sustainability footprint collection. A certain degree of noise can be added to the collected data to obscure individual data points but still allow for useful aggregate analysis.³

A classical challenge of such DP-based solutions would be to keep the utility (e.g., the statistical properties) of the sustainability data high for the system while still protecting the privacy of users and systems. In other words, the privacy budget—the amount of noise that can be added to the sustainability data without compromising privacy—needs to be determined by the sensitivity of the sustainability data being collected and the desired level of privacy protection. Another challenge for DP-based solutions is to keep the total noise added by all parties within an acceptable range, and failure to do so requires a trusted aggregator to correct the noise. Since DP-based solutions protect the data owner by providing dataset indistinguishability, they can be used as a privacy-preserving way of releasing data. However, one has to ensure their correct-by-construction nature^{8,11} while adopting them.

To provide cryptographic guarantees and to preserve the utility of sustainability data to a higher extent compared to DP, an alternate solution is to use homomorphic encryption (HE).⁴¹ With this solution, carbon sources can encrypt sustainability data as well as enable the decision-making agent to measure/compute any statistical information on those encrypted data.

There are, however, several challenges associated with this solution. HE requires significant computational resources and can increase the size of the actual data (because of encryption) being transmitted, making them more difficult to store and transmit efficiently. Furthermore, there are currently limitations on the types of computations that can be performed on homomorphically encrypted data. For example, HE schemes support only addition and multiplication. Complex operations, such as division or trigonometric functions, may not be efficiently supported.

While the direct use of HE may not be appropriate for resource-constrained carbon emission sources, further research is warranted to check whether optimized versions of HE, such as partial HE, leveled HE, and threshold HE, can be utilized or a new, lightweight, secure, and bespoke HE (e.g., selective HE) needs to be designed for sustainability in data centers. Nevertheless, many major chip/system vendors, such as Intel, AMD, and ARM, are actively exploring hardware support for HE, and when this is available, it will provide a trusted basis for implementing challenges to many of the security solutions identified in this article.

Another alternative approach involving less computational overhead than HE is zero-knowledge proofs,⁶ in which carbon sources can demonstrate to the sustainability certification agent that sustainability footprints are valid without disclosing the actual values that would otherwise compromise privacy.

However, zero-knowledge proofs can be used only to prove the authenticity of sustainability data and are not intended for analyzing and making any decisions. To address the challenges of each solution, further investigation is needed to determine whether HE or DP can be combined with zero-knowledge proofs.

Privacy-Preserving Footprint Aggregation

Collecting and processing sustainability data from multiple sites in data centers require secure collaboration among multiple untrusted parties, including cloud operators, regulators, and users, each with its own confidentiality, privacy, security, and trust requirements. While being aggregated either in centralized or distributed data centers, sustainability data can still reveal sensitive information about users and systems, as discussed in the “Security Challenges for Sustainability” section. Therefore, the high-level goals are to 1) perform aggregation, summary, or other functions on sustainability data whose results do not disclose information about the underlying data and 2) ensure that aggregations provide (provably) accurate higher-level data without exposing underlying sensitive information, e.g., proof of the sustainability compliance of a manufacturing process without exposing unit-wise behaviors or specific metrics.

Potential solutions. A plausible approach to privacy-preserving aggregation for sustainability data involves secure multiparty communication (MPC), in which multiple carbon footprint aggregators located at different locations collaborate to perform computations on their combined data without revealing any individual data points.⁵ MPC requires minimal trust and aims to ensure that each party’s input is kept private while allowing the parties to compute the desired aggregation, summary, or other functions on their combined data whose results do not disclose information about the underlying data. One such MPC platform is Confidential Space by Google (<https://cloud.google.com/blog/products/identity-security/announcing-confidential-space>), which allows sustainability data to be encrypted and stored in a TEE that only authorized workloads are allowed to access. Additionally, such data are isolated from other workloads running on the same machine, protecting against unauthorized access and tampering. MPC-based solutions, however, incur higher computational and communication overheads due to secure computations and sharing of encrypted results.

To minimize sustainability data movement, federated learning can be used, in which training a machine learning model (e.g., carbon footprint optimization) on decentralized sustainability data/metrics can be performed without having to transfer the data to a centralized location.

Each site of the distributed data center will train a local model on its sustainability data and send the updated model weights to a central server, which aggregates them to create a global model. This approach allows data to remain local and private while still benefiting from a centralized learning process. Note that existing federated learning techniques are susceptible to model poisoning and model stealing attacks; this further imposes challenges to adopt federated learning-based solutions for aggregating sustainability data.

Public Sustainability Ledgers

Public sustainability ledgers can be used for tracking carbon emissions or energy consumption and thus can provide transparency and accountability in the management of resources. However, there are also security and privacy issues that need to be considered when using these public ledgers. For example, if public ledgers contain sensitive data (e.g., carbon credit allocations, sales, and expenditures) about the sustainability practices of individuals and organizations, attackers may track the individuals/organizations or infer proprietary algorithms. Also, sustainability data may be stored on multiple public ledgers or private databases, which may not be interoperable. This can create challenges in ensuring data consistency and accuracy and may also lead to data breaches if not properly secured.

Potential solutions. In combination with privacy-preserving measures, such as HE, zero-knowledge proofs, multiparty computations, and DP, public ledgers for sustainability reporting can be provided through smart contracts¹ deployed on the public blockchain. The smart contract records the sustainability footprints from different sources and stores the encrypted records in blocks on the blockchain. The sustainability footprints submitted to the blockchain undergo verification by the participating entities through a consensus mechanism, such as proof of work or proof of stake. This ensures the accuracy and integrity of the recorded footprints. Consumers, stakeholders, and regulators can access the public blockchain to track and verify the provenance of sustainability footprints. Although smart contracts, in concert with a verifiable sustainability footprint collection architecture (Figure 1) and privacy-preserving measures, can offer secure and public sustainability ledgers, smart contracts can also be subject to vulnerabilities that can be exploited by attackers. Therefore, it is important to thoroughly test and audit smart contracts to ensure their security and reliability. Furthermore, blockchain technology can be used to address the inconsistency and data breach issues of distributed public ledgers. However, current blockchain technologies are susceptible to various types of attacks, including 51% (majority) attacks and DoS attacks. Thus, it is important to ensure that the blockchain network is properly secured and that appropriate security measures are in place to prevent such attacks.

While the potential security solutions outlined in this article may contribute to carbon footprints, future research is necessary to rigorously evaluate the performance and security guarantees of the existing and newly designed solutions. As discussed in the “Why Is Sustainability a Security Problem?” section, the importance of such security solutions in ensuring the trustworthiness of sustainability data and incentivizing users toward sustainability practices is crucial for addressing global climate change and is believed to outweigh the impact of systems lacking such guarantees.

Enhancing Standardization of Security Mechanisms

Security mechanisms are essential to ensure compliance with regulations and standards, preventing unauthorized access and the exposure, tampering, or misuse of sustainability data. Irrespective of the specific solution used to ensure the security of sustainability, a common need is to ease the adoption of those mechanisms and reduce their footprint, both in terms of performance and sustainability. For instance, a TEE-based solution for verifiable data collection or an HE-based approach for privacy-preserving footprint collection should be lightweight and have a small footprint so

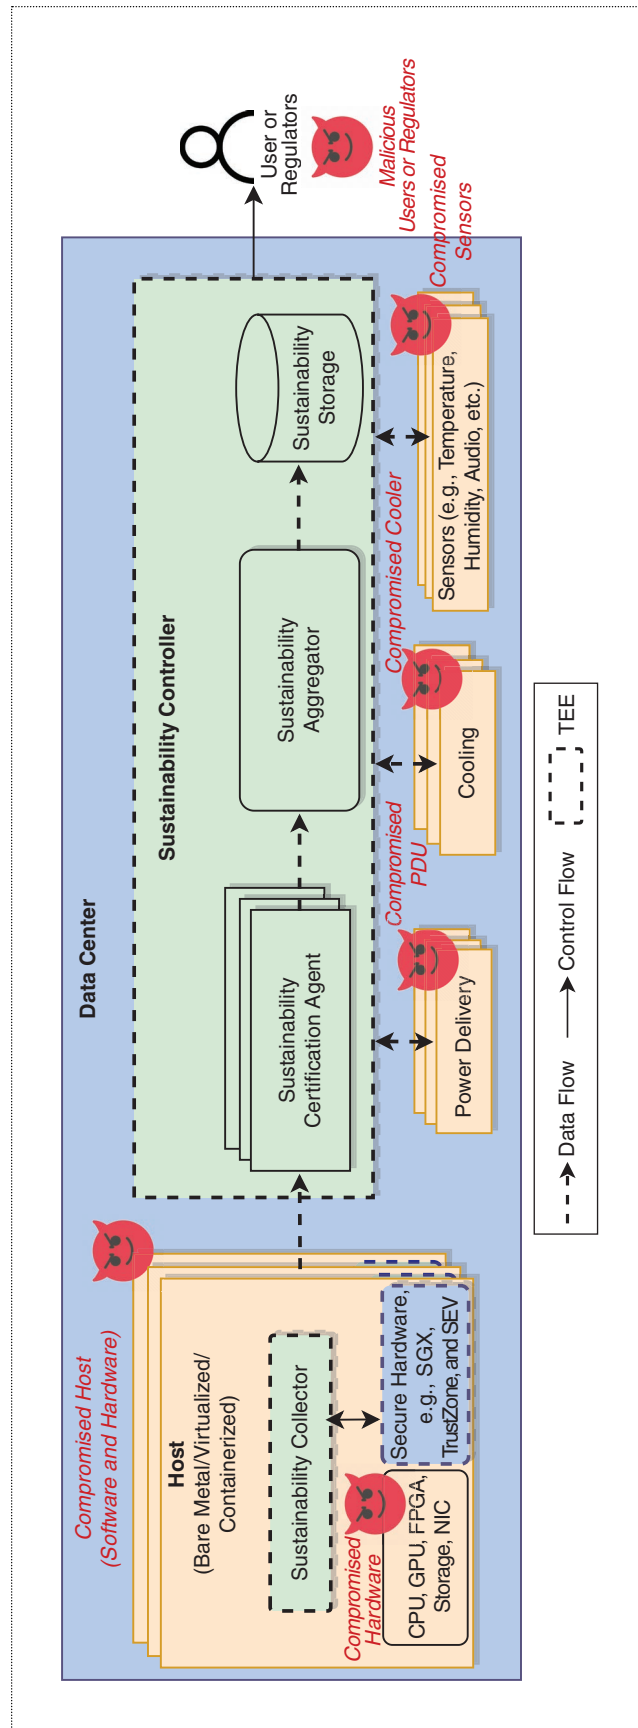


Figure 1. To enable the verifiability of sustainability metrics, we propose that sustainability-aware data centers be equipped with a sustainability collector, certification agent, sustainability aggregator, and sustainability storage. We mark components in the data center with an adversary symbol to denote potentially compromised components. Unchanged items in data centers are shaded blue, modified items are shaded orange, and new items added for sustainability are in green. FPGA: field-programmable gate array; NIC: network interface card.

as to minimize overall carbon consumption. As trustworthiness is foundational in sustainability initiatives, stakeholders, including governments, businesses, and users, need high security and privacy assurance of sustainability data, which is crucial for the success and adoption of sustainability practices. Since sustainability data are critical for understanding trends and for long-term planning and monitoring to counter global issues, such as climate change, it is necessary to rigorously evaluate the effectiveness of security measures toward sustainability to make informed decisions for a sustainable future. As sustainability is a global concern that requires collaboration across borders, standardizing security mechanisms for sustainability data will accelerate their adoption in other sectors, facilitate international cooperation, and ensure consistent protection standards and interoperability. Incentives and regulations need to be introduced to motivate organizations to adopt and implement standardized security mechanisms. These could include tax incentives, certification programs, or regulatory requirements that prioritize sustainability and security. Note that all challenges toward sustainability cannot be solved with technical solutions alone. Hence, offering both fundamental principles and secure guarantees is more likely to assist in the development of policies. This, in turn, can contribute to and accelerate the global effort to combat climate change. Without robust policies, all optimizations are susceptible to the Jevons paradox (i.e., increasing efficiency can lead to increased consumption), which signifies that both regulation and security are crucial components. Hence, collaboration and cooperation among industry players, researchers, and policymakers are necessary to establish these common goals and objectives.

Security infrastructure for a sustainable system is indispensable for protecting the environment and our planet. The central goal of this security infrastructure is to enable service providers to produce unforgeable proofs of sustainability footprints for users or regulators while preventing potential security and privacy threats by malicious users or compromised systems. Toward this goal, this article discussed the threat landscapes and new security challenges to achieve sustainability of data centers and presented potential research directions to develop primitives that allow domain experts to construct and operate sustainable data centers. The proposed challenges and potential solutions also lay the foundations for other sustainable systems, such as manufacturing, telecommunication systems, and automated transportation systems. ■

Acknowledgment

The work reported in this article has been supported by the National Science Foundation, under Grants 2215017, 2214980, 2046444, 2215016, 1750109, 1900706, 1918225, 1951880, 2106263, and 2106434.

References

1. M. Aloqaily, A. Boukerche, O. Bouachir, F. Khalid, and S. Jangsher, "An energy trade framework using smart contracts: Overview and challenges," *IEEE Netw.*, vol. 34, no. 4, pp. 119–125, Jul./Aug. 2020, doi: 10.1109/MNET.011.1900573.
2. C. C. Tsai, D. E. Porter, and M. Vij, "Graphene-SGX: A practical library OS for unmodified applications on SGX," in *Proc. USENIX Annu. Tech. Conf. (USENIX ATC)*, Santa Clara, CA, USA: USENIX Association, Jul. 2017, pp. 645–658.
3. C. Dwork, "Differential privacy," in *Proc. Automata, Lang. Program., 33rd Int. Colloq. (ICALP)*, Venice, Italy. Berlin, Germany: Springer-Verlag, Jul. 10–14, 2006, pp. 1–12, doi: 10.1007/11787006_1.
4. A. Gandhi et al., "Metrics for sustainability in data centers," *ACM SIGENERGY Energy Inform. Rev.*, vol. 3, no. 3, pp. 40–46, 2023.
5. O. Goldreich, "Secure multi-party computation," Weizmann Institute of Science, Rehovot, Israel, 1998. [Online]. Available: <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=dce0d462c182121f37279e3809d484624f3d3eba>
6. S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems," in *Proc. 17th Annu. ACM Symp. Theory Comput. (STOC)*, New York, NY, USA: Association for Computing Machinery, 1985, pp. 291–304, doi: 10.1145/22145.22178.
7. H. Hlavacs, T. Treutner, J.-P. Gelas, L. Lefevre, and A.-C. Orgerie, "Energy consumption side-channel attack at virtual machines in a cloud," in *Proc. IEEE 9th Int. Conf. Dependable, Autonomic Secure Comput.*, Piscataway, NJ, USA: IEEE Press, 2011, pp. 605–612, doi: 10.1109/DASC.2011.110.
8. J. Jin, E. McMurtry, B. I. Rubinstein, and O. Ohrimenko, "Are we there yet? Timing and floating-point attacks on differential privacy systems," in *Proc. IEEE Symp. Secur. Privacy (SP)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 473–488, doi: 10.1109/SP46214.2022.9833672.
9. E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, 2020, doi: 10.1126/science.aba3758.
10. E. McKenna, I. Richardson, and M. Thomson, "Smart meter data: Balancing consumer privacy concerns with legitimate applications," *Energy Policy*, vol. 41, pp. 807–814, Feb. 2012, doi: 10.1016/j.enpol.2011.11.049.
11. I. Mironov, "On significance of the least significant bits for differential privacy," in *Proc. ACM Conf. Comput.*

- Commun. Secur.*, New York, NY, USA: ACM, 2012, pp. 650–661, doi: 10.1145/2382196.2382264.
12. C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramian, “Carbon-aware energy capacity planning for datacenters,” in *Proc. 20th IEEE Int. Symp. Model. Anal. Simul. Comput. Telecommun. Syst. (MASCOTS)*, 2012, pp. 391–400, doi: 10.1109/MASCOTS.2012.51.
 13. T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, “Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds,” in *Proc. 16th ACM Conf. Comput. Commun. Secur. (CCS)*, New York, NY, USA: ACM, 2009, pp. 199–212, doi: 10.1145/1653662.1653687.
 14. I. Shumailov, Y. Zhao, D. Bates, N. Papernot, R. Mullins, and R. Anderson, “Sponge examples: Energy-latency attacks on neural networks,” in *Proc. IEEE Eur. Symp. Secur. Privacy (EuroS&P)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 212–231, doi: 10.1109/EuroSP51992.2021.00024.
 15. J. Zhu et al., “Enabling rack-scale confidential computing using heterogeneous trusted execution environment,” in *Proc. IEEE Symp. Secur. Privacy (SP)*, 2020, pp. 1450–1465, doi: 10.1109/SP40000.2020.00054.
 16. “Why data centres are the new frontier in the fight against climate change,” *Computerworld*, Aug. 2019. [Online]. Available: <https://www.computerworld.com/article/3431148/why-data-centres-are-the-new-frontier-in-the-fight-against-climate-change.html>
 17. M. A. B. Siddik, A. Shehabi, and L. Marston, “The environmental footprint of data centers in the United States,” *Environ. Res. Lett.*, vol. 16, no. 6, 2021, Art. no. 064017, doi: 10.1088/1748-9326/abfba1.
 18. “Achieving breakthrough results in sustainability.” Bain. [Online]. Available: <https://www.bain.com/insights/achieving-breakthrough-results-in-sustainability>
 19. “Volkswagen: The scandal explained,” *BBC News*, Dec. 2015. [Online]. Available: <https://www.bbc.com/news/business-34324772>
 20. “Power requirements to train modern LLMs.” nnlabs.org. [Online]. Available: <https://www.nnlabs.org/power-requirements-of-large-language-models/>
 21. D-REC. [Online]. Available: <https://dreco.org/>
 22. “Driving accountable sustainability in the consumer industry.” Deloitte. [Online]. Available: <https://www2.deloitte.com/us/en/insights/industry/retail-distribution/accountable-sustainability-consumer-industry.html>
 23. “NSF 20-594: NSF/VMware partnership on the next generation of sustainable digital infrastructure (NGSDI),” National Science Foundation, Alexandria, VA, USA, 2020. [Online]. Available: <https://www.nsf.gov/pubs/2020/nsf20594/nsf20594.htm>
 24. K. P. Pucker, “Overselling sustainability reporting,” *Harvard Bus. Rev.*, May/June 2021. [Online]. Available: <https://hbr.org/2021/05/overselling-sustainability-reporting>
 25. “How much energy do data centers really use?” Energy Innovation. [Online]. Available: <https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/>
 26. “CO₂e.” U.S. Environmental Protection Agency. [Online]. Available: <https://www3.epa.gov/carbon-footprint-calculator/tool/definitions/co2e.html>
 27. R. Miller. “The bitcoin energy debate: Lessons from the data center industry.” Data Center Frontier. [Online]. Available: <https://www.datacenterfrontier.com/energy/article/11428174/the-bitcoin-energy-debate-lessons-from-the-data-center-industry>
 28. “Carbon emissions of data usage increasing, but what is yours?” Climate Neutral Group. [Online]. Available: <https://www.climateneutralgroup.com/en/news/carbon-emissions-of-data-center>
 29. “AI power consumption exploding.” Semiconductor Engineering. [Online]. Available: <https://semiengineering.com/ai-power-consumption-exploding/>
 30. “Microsoft and 20 others call for carbon accounting standards.” DatacenterDynamics. [Online]. Available: <https://www.datacenterdynamics.com/en/news/microsoft-and-20-others-call-for-carbon-accounting-standards/>
 31. “Driverless cars: The dark side of autonomous vehicles that no one’s talking about.” Euronews. [Online]. Available: <https://www.euronews.com/green/2023/01/16/driverless-cars-the-dark-side-of-autonomous-vehicles-that-no-ones-talking-about>
 32. “Reading the runes: EU data center regulations are coming sooner than you think.” DatacenterDynamics. [Online]. Available: <https://www.datacenterdynamics.com/en/marketwatch/reading-the-runes-eu-data-center-regulations-are-coming-sooner-than-you-think/>
 33. “Federal sustainability plan: Catalyzing America’s clean energy industries and jobs,” The White House, Washington, DC, USA, Dec. 2021. [Online]. Available: <https://www.sustainability.gov/pdfs/federal-sustainability-plan.pdf>
 34. T. Donaldson, “There’s more cheating going on in sustainability than you think,” *Sourcing J.*, Oct. 2017. [Online]. Available: <https://sourcingjournal.com/topics/sustainability/apparel-sustainability-unethical-practices-73619/>
 35. J. Greig, “CISA, Clarity highlight severe vulnerabilities in popular power distribution unit product,” *The Record*, Sep. 2022. [Online]. Available: <https://therecord.media/cisa-clarity-highlight-severe-vulnerabilities-in-popular-power-distribution-unit-product>
 36. “Physical infrastructure cybersecurity: A growing problem for data centers.” Data Center Knowledge. [Online]. Available: <https://www.datacenterknowledge.com/security/physical-infrastructure-cybersecurity-growing-problem-data-centers>
 37. N. Nelson. “Bug in the Linux kernel allows privilege escalation, container escape.” Threatpost. [Online]. Available: <https://threatpost.com/bug-linux-kernel-privilege-escalation-container-escape/178808/>

38. D. Jacobs and L. P. Kalbers, "The Volkswagen diesel emissions scandal and accountability," *CPA J. Content*, Jul. 2019. [Online]. Available: <https://www.cpajournal.com/2019/07/22/9187/>
39. "Serious security: The Samba logon bug caused by outdated crypto," *Sophos News*, Jan. 2023. [Online]. Available: <https://news.sophos.com/en-us/2023/01/30/serious-security-the-samba-logon-bug-caused-by-outdated-crypto/>
40. Heartbleed Bug. [Online]. Available: <https://heartbleed.com/>
41. C. Gentry, "A fully homomorphic encryption scheme," Ph.D. Dissertation, Stanford University, Stanford, CA, USA, 2009.

Syed Rafiul Hussain is an assistant professor in the Department of Computer Science and Engineering, Pennsylvania State University, State College, PA 16802 USA. His research interests include systems and network security, formal methods, and sustainability. Hussain received a Ph.D. in computer science from Purdue University. He is a Member of IEEE and the Association for Computing Machinery. Contact him at hussain1@psu.edu.

Patrick McDaniel is the Tsun-Ming Shih Professor of Computer Sciences in the School of Computer, Data, and Information Sciences, University of Wisconsin-Madison, Madison, WI 53706 USA. His research interests include computer and network security and technical public policy, with a focus on mobile device security, the security of machine learning systems, program analysis for security, sustainability, and election systems. McDaniel received a Ph.D. in computer science and engineering from the University of Michigan. He is a recipient of the Association for Computing Machinery (ACM) Special Interest Group on Operating Systems Hall of Fame Award and ACM Special Interest Group on Security, Audit, and Control Outstanding Innovation Award. He is a Fellow of IEEE, the ACM, and the American Association for the Advancement of Science. Contact him at mcdaniel@cs.wisc.edu.

Anshul Gandhi is an associate professor in the Department of Computer Science, Stony Brook University, Stony Brook, NY 11790 USA. His research interests include performance analysis, modeling, and evaluation; distributed systems; and sustainability. Gandhi received a Ph.D. in computer science from Carnegie Mellon University. He is a Senior Member of IEEE and the Association for Computing Machinery. Contact him at anshul@cs.stonybrook.edu.

Kanad Ghose is a State University of New York Distinguished Professor of Computer Science at Binghamton University, Binghamton, NY 13902 USA. His research interests include energy-aware systems at all scales, processor microarchitectures, and hardware security. Ghose received a Ph.D. in computer science from Iowa State University. He is a Member of IEEE and the Association for Computing Machinery. Contact him at ghose@binghamton.edu.

Kartik Gopalan is a professor in the Department of Computer Science, Binghamton University, Binghamton, NY 13902 USA. His research interests include computer systems, including virtualization, security, operating systems, networks, and sustainability. Gopalan received a Ph.D. in computer science from Stony Brook University. He is a Senior Member of IEEE and a member of the Association for Computing Machinery. Contact him at kartik@binghamton.edu.

Dongyoon Lee is an assistant professor in the Department of Computer Science, Stony Brook University, Stony Brook, NY 11790 USA. His research interests include compilers, operating systems, computer architecture, security, and sustainability. Lee received a Ph.D. in computer science engineering from the University of Michigan. He is a Member of IEEE and the Association for Computing Machinery. Contact him at dongyoon@cs.stonybrook.edu.

Yu David Liu is a professor in the Department of Computer Science, Binghamton University, Binghamton, NY 13902 USA. His research interests include programming languages, software engineering, formal methods for security, and sustainable and energy-aware applications and systems. Liu received a Ph.D. in computer science from Johns Hopkins University. He is a member of the Association for Computing Machinery. Contact him at davidl@binghamton.edu.

Zhenhua Liu is an associate professor in the Department of Applied Mathematics and Statistics and the Department of Computer Science, Stony Brook University, Stony Brook, NY 11790 USA. His research interests include optimization, machine learning, big data systems, and sustainable and energy-aware applications and systems. Liu received a Ph.D. in computer science from the California Institute of Technology. He is a member of the Association for Computing Machinery. Contact him at zhenhua.liu@stonybrook.edu.

Shuai Mu is an assistant professor in the Department of Computer Science, Stony Brook University, Stony Brook, NY 11790 USA. His research interests include computer systems, including virtualization, security, operating systems, networks, and sustainability. Mu received a Ph.D. in computer science from Stony Brook University. He is a Member of IEEE and the Association for Computing Machinery. Contact him at shuai.mu@stonybrook.edu.

Brook, NY 11790 USA. His research interests include distributed systems, multicore systems, and sustainable systems. Mu received a Ph.D. in computer science from Tsinghua University. He is a member of the Association for Computing Machinery. Contact him at shuai@cs.stonybrook.edu.

Erez Zadok is a professor at Stony Brook University, Stony Brook, NY 11790 USA. His research interests include

computer systems, storage systems, security, performance optimizations, and sustainability. Zadok received a Ph.D. in computer science from Columbia University. He is a senior member of the IEEE Computer Society, a distinguished member of the Association for Computing Machinery, and a member of USENIX. Contact him at erez.zadok@stonybrook.edu.