Security & Privacy Magazine

Verifiable Sustainability in Data Centers

Syed Rafiul Hussain, *Pennsylvania State University, University Park, Pennsylvania, 16802, USA*
Patrick McDaniel, *University of Wisconsin–Madison, Madison, Wisconsin, 53706, USA*
Anshul Gandhi, *Stony Brook University, Stony Brook, New York, 11790, USA*
Kanad Ghose, *Binghamton University, Binghamton, New York, 13902, USA*
Kartik Gopalan, *Binghamton University, Binghamton, New York, 13902, USA*
Dongyoon Lee, *Stony Brook University, Stony Brook, New York, 11790, USA*
Yu David Liu, *Binghamton University, Binghamton, New York, 13902, USA*
Zhenhua Liu, *Stony Brook University, Stony Brook, New York, 11790, USA*
Shuai Mu, *Stony Brook University, Stony Brook, New York, 11790, USA*
Erez Zadok, *Stony Brook University, Stony Brook, New York, 11790, USA*

Abstract—Sustainability is crucial for combating climate change and protecting our planet. While there are various systems that can pose a threat to sustainability, data centers are particularly significant due to their substantial energy consumption and environmental impact. Although data centers are becoming increasingly accountable to be sustainable, the current practice of reporting sustainability data is often mired with simple green-washing. To improve this status quo, users as well as regulators need to verify the data on the sustainability impact reported by data center operators. To do so, data centers must have appropriate infrastructures in place that provide the guarantee that the data on sustainability is collected, stored, aggregated, and converted to metrics in a secure, unforgeable, and privacy-preserving manner. Therefore, this paper first introduces the new security challenges related to such infrastructure, how it affects operators and users, and potential solutions and research directions for addressing the challenges for data centers and other industry segments.

Sustainability is the practice of performing human activities in ways that do not leave lasting harmful effects [58]. Unfortunately, the harm to the planet is clearly growing, 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 [25]. 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, desertization, and loss of biodiversity. Failure to address this lack of sustainability now will create long-term problems for future generations [60].

Today, achieving the goals of sustainability requires the honest, best efforts of humans and an 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 Environmental Protection Agency (EPA) about emissions when under test [39]. 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

XXXX-XXX © 2021 IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

the alarming rise of computation and the pervasive use of artificial intelligence (*e.g.*, ChatGPT [6]), data centers pose many negative impacts on the environment caused by energy use, hardware manufacturing and disposal, building maintenance, and other factors. Indeed, a recent study showed that over 2–4% of all energy used worldwide was by data centers [53], [43]. 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 from green generation sources [1] or by paying other entities to be sustainable. This signifies a lack of transparency and accountability that hinder 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.

The lack of accountability and transparency to address sustainability is primarily rooted in the absence of *complete* and *verifiable* sustainability data and metrics [32], [18], [48]. Comprehensive and fine-grained sustainability metrics [3] 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 reliable audit trail of carbon consumption, ensure accurate and precise accountability and compliance benefits (*e.g.*, accurately identify entities who made changes or performed certain actions), and optimize system performance [12], [38], [3]. 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 lifecycle 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 [4]. For example, there are dozens of different ways to calculate data on global data center energy consumption based on public and private data—each resulting in an assessment that is often contradictory with others [63]. 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, exposure of sensitive users' 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 paper, the infrastructures used to measure and maintain *operational sustainability* (*i.e.*, environmental footprints transpired 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 data and the enforcement of corresponding security and privacy policies. Hence, it is imperative that the environmental footprint caused by data center operations can be verified by interested third parties (*e.g.*, the EPA [73], citizen scientists, and the public).

This article, therefore, looks at the security issues in the sustainability data pipeline comprising 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.

1. 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, 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. Therefore, in this article we specifically focus on sustainability in data centers.

Operations within data centers already contribute significantly to the global carbon footprint [67], [56]. The rise in popularity of resource-intensive Big Data, AI, cryptocurrency, and Machine-Learning workloads is poised to make data center operations even more unsustainable [42], [10], [53], [69], [55]. 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 [22]. 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 [76], [6], [11], and applications based on LLMs, imposing a much heavier toll on the environment.

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 [1]. Similar greenwashing practices have also been observed in other sectors such as autonomous vehicles [30] and telecommunication industries [23]. Such deceptive approaches undermine the transparency and credibility of sustainability claims, making it difficult for stakeholders to make informed decisions. The European Union's Corporate Sustainability Reporting Directive (CSRD) [15] mandates that by 2024, corporations have to report non-financial 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 [16], [24]. In the U.S., we also see the beginnings of directives similar to the EUs [71], but details are still emerging.

2. Why is Sustainability a Security Problem?

Ensuring the accuracy and credibility of sustainability metrics, as well as supporting audits, 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 [73], [43], sustainability metrics in data centers require the authenticity, confidentiality, integrity, and availability of data collected, processed, stored, and used locally within a data center [34]. However, unlike traditional cloud computing systems where the focus is primarily on security and privacy of user applications and data [17], [79], [20], 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 induced within a data center to preserve the privacy of users' or operators' sustainability data. It is thus imperative to ensure the security of (i) data collection processes, (ii) the process of generating verifiable, easily auditable sustainability metrics, and (iii) 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.

2.1. Threat Models for Sustainability in Data Centers

Although the trust assumptions and threat models for sustainable systems may vary widely based on the system architecture and requirements, the threat models for a sustainable data center can be primarily derived with respect to three entities: (a) the service provider, (b) the users, and (c) third-party observers (*e.g.*, regulatory agencies). One may assume that the service provider can be considered to be trusted but the underlying infrastructure (*e.g.*, OS and services) provided by third-party vendors/suppliers can be untrusted or become compromised, whereas others may assume that both the service provider and the underlying infrastructure become rogue. For example, benign and unsuspecting data center providers often use virtual machines (VMs) or containers that are already offered by third-party infrastructure providers and can be loaded with backdoors or malware. A malicious infrastructure provider can deliberately manipulate energy consumption metrics, bypass sustainability regulations, and overcharge the data center provider for the total energy consumption. This not only undermines the data center provider's sustainability efforts but also leads to inflated costs and financial losses. Moreover, when users' jobs run in an environment where data center and/or infrastructure providers are malicious, attackers can gain unauthorized access to read or modify the job's code and data. For example, attackers may introduce unaccounted read/write operations [19], [52] to users' jobs which in turn inflate users' carbon footprints, leading to overbilling the customers and increasing the financial profits of data center and infrastructure providers. Such carbon footprint inflation can also be achieved by violating the integrity of the sustainability metrics (*e.g.*, code or data) [19], [52] or by manipulating the system traces and logs—the evidence trail of carbon consumption [45] by the compromised VMs or malicious processes in data centers. Similarly, compromised data center providers may report false carbon footprints to the regulators [14] to evade high CO2 taxes or regulations.

The other key entities in data centers (*i.e.*, users) can also be assumed to be untrusted as they may try to launch attacks (*e.g.*, DoS) against other users or data center providers, or obtain higher levels of service than they are allocated, and thus mislead the cloud service providers about the user's carbon usage. Last but not least, third-party observers (*e.g.*, regulatory agencies) may be tasked with verifying the footprint reported by the service provider in the process of executing policy or oversight (*e.g.*, by comparing sustainability costs reported by cloud operators, users, and utilities); but even these observers may be considered untrusted, as they could collude with others to mislead reporting, may have rogue insider elements within the data center, and may even be under political or other pressure to "fudge" or misrepresent the data.

2.2. New Security Challenges for Sustainability

Due to the complex design of data centers, which relies on intricate trust assumptions among numerous stakeholders, it is necessary to address diverse security threats ranging from malicious software/service providers (in Software as a Service or SaaS models), compromised operating systems or hypervisors (in Platform as a Service or PaaS models), or compromised sensors, devices, and firmware owned by infrastructure providers (in Infrastructure as a Service or IaaS models) to malicious or honest-but-curious users. In light of the above discussion, we discuss next some security challenges for a system aiming for sustainability and summarise those in Table 1. Note that the nature of threats will be different for different sustainable systems (*e.g.*, transportation, manufacturing) based on trust assumptions.

❑ **Lack of authenticity of carbon emission sources (C1).** Sensors and devices (*e.g.*, PDUs) reporting and computing sustainability data can be malicious and may become compromised due to unintentional vulnerabilities or intended backdoors in their hardware, firmware, and software [37]. As a result, by taking control of those sensors and devices, attackers may violate the authenticity and forge carbon footprints to launch nefarious attacks. For instance, attackers may cause over/under-billing to customers by forging/manipulating carbon consumption records. Attackers may also induce carbon-exhaustion attacks on other users by misreporting over-consumption of carbon, or evade compliance checking of regulatory agencies by misreporting low carbon emissions when operating in test mode (similar to Volkswagen's scandal [39]). 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 gleaning carbon footprints from multiple such sensors.

❑ **Untrustworthy physical environment (C2).** Sensors and apparatuses used to collect carbon footprint data can be subjected to direct and indirect data manipulation attacks. For example, an attacker having direct physical access to sensors or data structure infrastructure can manipulate sensors' readings to generate false sustainability data or manipulate the cooling system to disrupt sustainability operations [47]. Conversely, in indirect attacks, the attackers do not have direct physical access to sensors but exploit physical side-channels [26] between different components/sensors in data centers to affect sustainability operations and cause reputation loss to their competitors. 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.

❑ **Lack of access control and information flow control (C3).** Sharing physical resources such as hardware and sensors among multiple users introduces new challenge of isolating each tenant's data and ensuring that one tenant cannot access another's sustainability footprints. The lack of granular and dynamic access control configurations, and adequate resource isolation, can lead to the failure to ensure that each workload and its associated users have the appropriate access privileges to sustainability footprints, without compromising data security. Without proper access control and information flow-control measures, there is, therefore, a risk of unauthorized access to sensitive sustainability data, potentially leading to data breaches, privacy violations, and other security issues. Furthermore, sustainability data obtained from various sources can be illegitimately tampered with by malicious users processes or compromised system processes. 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 [59]. The lack of access control mechanisms, such as Discretionary Access Control (DAC), Mandatory Access Control (MAC), or combinations thereof, therefore, may enable attackers to manipulate (*i.e.*, add, modify, or remove) carbon footprint and sustainability states. As a result, the regular sustainability operations of the system are likely to be disrupted, which may cause the system to produce unwarranted carbon footprints or eliminate them. Tampering with sustainability data by adversaries (*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.

❑ **Sensitive information disclosure (C4).** Collecting sustainability data from disparate carbon sources (*e.g.*, sensors

ID	Vulnerabilities, Threats, and New Security Challenges	Impacts	Possible Ideas to Solutions
C1	Lack of authenticity of carbon emission sources allows malicious processes to forge, tamper, or misreport carbon usage	Cause over-/under-billing to customers by tampering with carbon usage, evade regulatory agencies by misreporting low carbon emissions	Verifiable footprint collection (§3.1)
C2	Untrustworthy physical environment may allow attackers to manipulate sensors and apparatuses within a data center directly or indirectly	Induce higher operational costs, cause over-/under-billing to customers, and denial-of-service attacks	Verifiable footprint collection (§3.1)
C3	Cryptographic flaws may allow forging the proof of carbon usage	Financial loss and disruption the data center operations	Verifiable footprint collection (§3.2)
C4 & C6	Disclosure of sustainability metrics to malicious service providers and other users 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 (§3.2, §3.3, & §3.4)
C5	Lack of or flaws in the access control or information flow control mechanisms may allow malicious processes (controlled by malicious users or service providers) to access and tamper with the databases storing carbon footprint trails	Exposure of users' private data such as location, behavior, and intellectual properties	Verifiable carbon footprint collection (§3.2)
C7	Evasive carbon offset techniques allow corporations 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 (§3.2)
C8	Multiple parties may collude to misreport carbon usage	Tax evasion, financial loss, and environmental hazards	Verifiable footprint collection (§3.2)

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

and PDUs) in an unregulated manner may disclose the sustainability metrics to service providers and other users. Such unauthorized exposure of footprint data will violate the privacy of user's data, location, behavior, and intellectual properties such as proprietary scheduling techniques, factors used for competitive pricing for service classes. Unauthorized access to footprint data can enable an adversary to prevent a co-tenant from realizing an improved sustainability target or even allow them to initiate DoS attacks on the co-tenant.

❑ **Cryptographic flaws (C5).** The ability of a sustainable system to provide proof of carbon footprint to users and regulators is essential for ensuring the trustworthiness of the system. Such proof of footprint should be built with cryptographic constructs. But flaws in the integration of cryptographic constructs with complex data center systems (*e.g.*, using weak cipher suites [21], [27]) or flaws in the implementations [2] 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, co-located servers, and poor isolation between different processes running on the same hardware in data centers, side-channel vulnerabilities (*e.g.*, page faults [75], cache misses [74], power [65] and timing [40] channels) may allow a malicious application to observe or tamper with carbon footprint patterns of other users' jobs/applications running on the same hardware. Such side-channels not only allow an attacker to fingerprint the data traffic of other users but also to extract the cryptographic keys or other confidential information of a user application by looking at the use of sustainability metrics [68]. Attackers 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).

❑ **Evasive carbon offset techniques (C7).** 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) [31]. This technique, also called carbon credit or climate credit, has been in practice for decades. It is often exploited by large corporations as it is extremely difficult, if not impossible, to track and verify if the amount of emissions balances out the amount of reductions [66], [70], [64]. 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 [1]. Similarly, Power Purchase Agreements (PPAs) [5] are used to have the data center operator finance the installation of a green energy producing farm, run, 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 REC and PPAs, the authenticity of green energy is, however, often kept out of sight of the users. The lack of authentication, therefore, enables corporations to make false claims about the energy source, while appearing in public to support sustainability efforts.

❑ **Collusion for evasion (C8).** Infrastructure providers and Power Distribution Unit (PDU) providers 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.

3. Research Directions for Countering Security Challenges for Sustainable Data Centers

Although many solutions [13], [78] have been designed for data-center security, most of them are not directly applicable to counter the security and privacy challenges towards sustainability as discussed in Section 2. 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.

3.1. 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 *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 to secure the system's carbon footprint measurement components. Such a trusted path should extend from the hardware level 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 tamper-proof 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 [8], Intel SGX [54], AMD SEV [44], and Keystone [49]. 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 [19], [9], [41]. These environments provide isolation guarantees needed to certify that metric data is collected and reported accurately, even in the presence of malicious applications, OS, or hypervisor. 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 within TEEs may introduce new security challenges to TEE-based solutions. Therefore, it is crucial to ensure high security assurance of these components through formal analysis. 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 the 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 SmartNICs, that can collect sustainability metrics from outside the host. For example, AWS Nitro [51] 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.

3.2. Privacy-Preserving Footprint Collection

Sustainability data collected through disparate carbon sources, such as sensors and PDUs in an unregulated manner, may incur unintended disclosure of sensitive data. Such exposure of footprint records would otherwise break the users' privacy, data, location, behavior, and intellectual properties such as proprietary scheduling techniques, trained machine learning models, and factors used for competitive

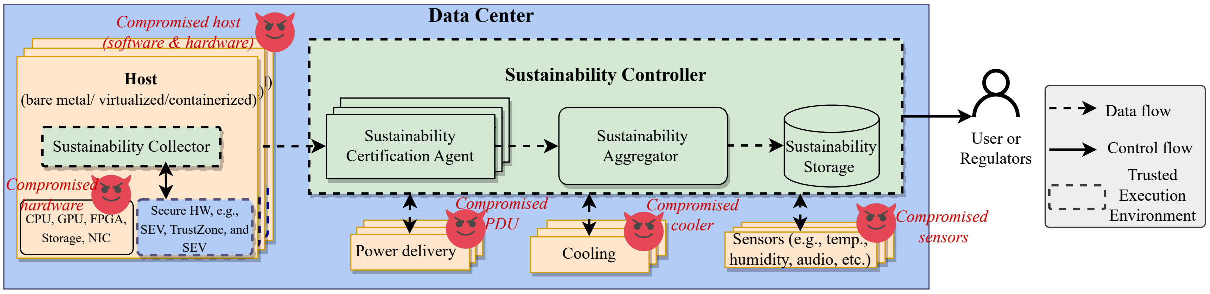
Data Center
Compromised host (software & hardware)
Host
(bare metal/ virtualized/containerized)
Sustainability Collector
Compromised hardware
CPU, GPU, FPGA, Storage, NIC
Secure HW, e.g., SEV, TrustZone, and SEV
Sustainability Controller
Sustainability Certification Agent
Sustainability Aggregator
Sustainability Storage
Compromised PDU
Power delivery
Compromised cooler
Cooling
Sensors (e.g., temp., humidity, audio, etc.)
Compromised sensors
User or Regulators
Data flow
Control flow
Trusted Execution Environment

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 potential compromised components. Unchanged items in data centers are shaded blue, modified items are shaded orange, and new items added for sustainability are in green.

pricing for service classes. Also, attackers may attempt to tamper with sensor data before it is 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 can be used as a plausible 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 [28], [29]. A classic challenge of such differential privacy-based solutions would be to keep the utility (*e.g.*, the statistical properties) of the data high to 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 potential solution is to use homomorphic encryption [35] that will allow the carbon sources to encrypt the sustainability data and enable the decision-making agent to measure/compute any statistical information on those encrypted data. There are, however, several challenges associated with this solution [57]. Homomorphic encryption requires significant computational resources and can increase the size of the actual data (because of encryption) being transmitted [57], making it more difficult to store and transmit efficiently. Furthermore, there are currently limitations [57] on the types of computations that can be performed on homomorphically encrypted data. For example, homomorphic encryption schemes support only addition and multiplication. Complex operations, such as division or trigonometric functions, may not be efficiently supported. This can limit the usefulness of this technique for some applications.

Another alternative approach would be to use zero-knowledge proofs [33], in which the 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 introduce new challenges as they induce high computational overhead.

3.3. Privacy-Preserving Footprint Aggregation

Collecting and processing sustainability data from multiple sites in data centers require secure collaboration between multiple untrusting parties, including cloud operators, regulators, and users, each with their 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 Section 2.2. Therefore, the high-level goals are to (1) perform aggregation, summary, or other functions on the 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 compliance of the manufacturing process without exposing unit-wise behaviors or specific metrics.

Potential Solutions: A plausible approach to privacy-preserving aggregation for sustainability data is secure multi-party communication (MPC) in which multiple parties collaborate to perform computations on their combined data without revealing any individual data points [36]. MPC aims to ensure each party's input is kept private while allowing them 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 [46], which would allow sustainability data to be encrypted and stored in a TEE that only authorized workloads are allowed to access. Additionally, such data is isolated from other workloads running on the same machine, providing protection against unauthorized access and tampering. Alternatively, federated learning [50] 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 [50].

3.4. 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 homomorphic encryption, zero-knowledge proofs, multi-party computations, and differential privacy, public ledgers for sustainability reporting can be provided through smart contracts [7] 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 (PoW) or Proof-of-Stake (PoS). 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 [61]. As such, it is important to thoroughly test and audit smart contracts to ensure their security and reliability [72]. Furthermore, blockchain technology [62] 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 denial-of-service attacks [77]. As such, it is important to ensure that the blockchain network is properly secured and that appropriate security measures are in place to prevent such attacks.

4. Enhancing Adoption and Standardization of Security Mechanisms

Irrespective of the specific solution used for the security artifacts, 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 trusted execution environment (TEE) based solution for verifiable data collection or a homomorphic encryption-based approach for privacy-preserving footprint collection should be lightweight and have small footprints so as to minimize overall carbon consumption. Standardizing the security mechanisms will also go a long way in accelerating their adoption in other sectors. Also, to raise awareness among stakeholders about the importance of security mechanisms for sustainable systems, and to educate them about the potential benefits (*e.g.*, data protection, compliance with regulations, reputation, and trust), workshops and seminars should be organized which will further improve user adoption and collaboration with other stakeholders. It is also necessary to educate the stakeholders about the challenges such as complexity and cost implications, evolving threats, and interoperability issues in adopting security solutions for sustainability. Industry standards and best practices need to be developed for security mechanisms in sustainable systems. These standards should address performance, interoperability, scalability, and energy efficiency aspects to ensure widespread adoption. 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. Overall, collaboration and cooperation among industry players, researchers, and policymakers are necessary to establish these common goals and objectives.

5. Conclusion

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. Towards this goal, this paper discusses the threat landscapes and new security challenges to achieve sustainability of data centers and presents 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.

Acknowledgement

The work reported in this paper has been supported by NSF under grants 2215017, 2214980, 2046444, 2215016, and 2106263.

6. REFERENCES

1. D-REC initiative. https://drecs.org/, April 2023.
2. The heartbleed bug. https://heartbleed.com/, 2023.
3. NSF/VMware partnership on the next generation of sustainable digital infrastructure (NGSDI). https://www.nsf.gov/pubs/2020/nsf20594/nsf20594.htm, 2023.
4. Overselling sustainability reporting. https://hbr.org/2021/05/overselling-sustainability-reporting, April 2023.
5. Power purchase aggrements. https://betterbuildingssolutioncenter.energy.gov/financing-navigator/option/power-purchase-agreement, April 2023.
6. AI. The power requirements to train modern large language models. https://www.nnlabs.org/power-requirements-of-large-language-models/, 2023.
7. Moayad Aloqaily, Azzedine Boukerche, Ouns Bouachir, Fariea Khalid, and Sobia Jangsher. An energy trade framework using smart contracts: Overview and challenges. *IEEE Network*, 34(4):119–125, 2020.
8. Arm Limited. Arm TrustZone Technology for the Armv8-M Architecture. https://developer.arm.com/documentation/100690/0201, 2016.
9. Sergei Arnautov, Bohdan Trach, Franz Gregor, Thomas Knauth, Andre Martin, Christian Priebe, Joshua Lind, Divya Muthukumaran, Dan O'Keeffe, Mark L. Stillwell, David Goltzsche, Dave Eyers, Rüdiger Kapitza, Peter Pietzuch, and Christof Fetzer. SCONE: Secure linux containers with intel SGX. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 689–703, Savannah, GA, November 2016. USENIX Association.
10. Esmail Asyabi, Azer Bestavros, Erfan Sharafzadeh, and Timothy Zhu. Peafowl: In-application CPU Scheduling to Reduce Power Consumption of In-memory Key-Value Stores. In *Proceedings of the 11th ACM Symposium on Cloud Computing*, SOCC '20, 2020.
11. Brian Bailey. AI power consumption exploding. https://semiengineering.com/ai-power-consumption-exploding/, 2022.
12. Noman Bashir, Tian Guo, Mohammad Hajiesmaili, David Irwin, Prashant Shenoy, Ramesh Sitaraman, Abel Souza, and Adam Wierman. Enabling sustainable clouds: The case for virtualizing the energy system. In *Proceedings of the ACM Symposium on Cloud Computing*, SoCC '21, pages 350–358, Seattle, WA, USA, 2021.
13. Stefan Berger, Ramón Cáceres, Dimitrios Pendarakis, Reiner Sailer, Enriquillo Valdez, Ronald Perez, Wayne Schildhauer, and Deepa Srinivasan. TVDc: Managing security in the trusted virtual datacenter. *ACM SIGOPS Operating Systems Review*, 42(1):40–47, 2008.
14. Eleanor Birrell, Anders Gjerdrum, Robbert van Renesse, Håvard Johansen, Dag Johansen, and Fred B. Schneider. SGX enforcement of use-based privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*, WPES'18, pages 155–167, New York, NY, USA, 2018. Association for Computing Machinery.
15. European Union Parliament Legislative Body. Legislative train 12.2022: Corporate sustainability reporting directive (CSRD). https://www.europarl.europa.eu/legislative-train/carriage/review-of-the-non-financial-reporting-directive/report?sid=6501, 2022.
16. G. Burton. Reading the runes: Eu data center regulations are coming sooner than you think. https://www.datacenterdynamics.com/en/marketwatch/reading-the-runes-eu-data-center-regulations-are-coming-sooner-than-you-think/, 2021.
17. Sean Carlin and Kevin Curran. *Cloud computing security*. IGI Global, 2013.
18. James Cascone, Justin Cook, Suzy O'Mara, Barb Renner, and Anthony Waelter. Driving accountable sustainability in the consumer industry. https://www2.deloitte.com/us/en/insights/industry/retail-distribution/accountable-sustainability-consumer-industry.html, April 2023.
19. Chia che Tsai, Donald E. Porter, and Mona Vij. Graphene-SGX: A practical library OS for unmodified

applications on SGX. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, pages 645–658, Santa Clara, CA, July 2017. USENIX Association.

20. Yanpei Chen, Vern Paxson, and Randy H Katz. What's new about cloud computing security. *University of California, Berkeley Report No. UCB/EECS-2010-5 January*, 20(2010):2010–5, 2010.
21. Thomas Claburn. The register: Fyi: Microsoft office 365 message encryption relies on insecure block cipher. https://threatpost.com/bug-linux-kernel-privilege-escalation-container-escape/178808/, 2022.
22. Climate Neutral Group. Carbon emissions of data usage increasing, but what is yours? https://www.climateneutralgroup.com/en/news/carbon-emissions-of-data-center, 2018.
23. Claire Curran. What will 5g mean for the environment? https://jsis.washington.edu/news/what-will-5g-mean-for-the-environment/, 2020.
24. J. Davis. Bring on regulations for data center sustainability, say Europe and APAC. https://journal.uptimeinstitute.com/bring-on-regulations-for-data-center-sustainability-say-europe-and-apac/, 2022.
25. Jenny Davis-Peccoud, Paul Stone, and Clare Tovey. Achieving breakthrough results in sustainability: Ceos who are passionate about change need to support the front line. https://www.bain.com/insights/achieving-breakthrough-results-in-sustainability, 2016.
26. Wenbo Ding, Hongxin Hu, and Long Cheng. Iotsafe: Enforcing safety and security policy with real iot physical interaction discovery. In *the 28th Network and Distributed System Security Symposium (NDSS 2021)*, 2021.
27. Paul Ducklin. Naked security: Serious security: The samba logon bug caused by outdated crypto. https://nakedsecurity.sophos.com/2023/01/30/serious-security-the-samba-logon-bug-caused-by-outdated-crypto/, 2023.
28. Cynthia Dwork. Differential privacy. In *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II 33*, pages 1–12. Springer, 2006.
29. Cynthia Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation: 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings 5*, pages 1–19. Springer, 2008.
30. Charlotte Elton. Driverless cars: The dark side of autonomous vehicles that no one's talking about. https://www.euronews.com/green/2023/01/16/driverless-cars-the-dark-side-of-autonomous-vehicles-that-no-ones-talking-about, 2020.
31. William Farrington. Tesla continues to cash in on carbon credits. https://www.proactiveinvestors.com/companies/news/988168/tesla-continues-to-cash-in-on-carbon-credits-988168.html, November 2022.
32. Caio Ferreira, Fabio Natalucci, Ranjit Singh, and Felix Suntheim. How strengthening standards for data and disclosure can make for a greener future. https://www.imf.org/en/Blogs/Articles/2021/05/13/how-strengthening-standards-for-data-and-disclosure-can-make-for-a-greener-future, April 2023.
33. Uriel Fiege, Amos Fiat, and Adi Shamir. Zero knowledge proofs of identity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 210–217, 1987.
34. Anshul Gandhi, Kanad Ghose, Kartik Gopalan, Syed Rafiul Hussain, Dongyoon Lee, David Liu, Zhenhua Liu, Patrick McDaniel, Shuai Mu, and Erez Zadok10. Metrics for sustainability in data centers. In *Proceedings of the 1st Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon'22)*, 2022.
35. Craig Gentry. *A fully homomorphic encryption scheme*. Stanford university, 2009.
36. Oded Goldreich. Secure multi-party computation. *Manuscript. Preliminary version*, 78(110), 1998.
37. Jonathan Greig. CISA, Claroty highlight severe vulnerabilities in popular power distribution unit product. https://therecord.media/cisa-claroty-highlight-severe-vulnerabilities-in-popular-power-distribution-unit-product, 2022.
38. Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. Chasing carbon: The elusive environmental footprint of computing. *IEEE Micro*, 42(4):37—47, jul 2022.
39. Russell Hotten. Volkswagen: The scandal explained. *British Broadcast Company*, dec 2015.
40. Ralf Hund, Carsten Willems, and Thorsten Holz. Practical timing side channel attacks against kernel space aslr. In *2013 IEEE Symposium on Security and Privacy*, pages 191–205. IEEE, 2013.
41. Mohit Kumar Jangid, Guoxing Chen, Yinqian Zhang, and Zhiqiang Lin. Towards formal verification of state continuity for enclave programs. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 573–590. USENIX Association, August 2021.
42. Kostis Kaffes, Dragos Sbirlea, Yiyan Lin, David Lo, and Christos Kozyrakis. Leveraging Application Classes to Save Power in Highly-Utilized Data Centers. In *Proceedings of the 11th ACM Symposium on Cloud Computing*, SOCC '20, 2020.
43. George Kamiya. Data centres and data transmission

networks. https://www.iea.org/reports/data-centres-and-data-transmission-networks, 2022.

44. David Kaplan, Jeremy Powell, and Tom Woller. Introduction to Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV). Technical report, Advanced Micro Devices, Inc., April 2016.
45. Vishal Karande, Erick Bauman, Zhiqiang Lin, and Latifur Khan. SGX-Log: securing system logs with SGX. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, ASIA CCS '17, pages 19–30, New York, NY, USA, 2017. Association for Computing Machinery.
46. Rene Kolga and Nelly Porter. Introducing confidential space to help unlock the value of secure data collaboration. https://cloud.google.com/blog/products/identity-security/announcing-confidential-space, October 2022.
47. Maria Korolov. Physical infrastructure cybersecurity: A growing problem for data centers. https://www.datacenterknowledge.com/security/physical-infrastructure-cybersecurity-growing-problem-data-centers, 2023.
48. Addisu Lashitew. The risks of us-eu divergence on corporate sustainability disclosure. https://www.brookings.edu/blog/future-development/2021/09/28/the-risks-of-us-eu-divergence-on-corporate-sustainability-disclosure/, 2023.
49. Dayeol Lee, David Kohlbrenner, Shweta Shinde, Krste Asanović, and Dawn Song. Keystone: An open framework for architecting trusted execution environments. In *Proceedings of the Fifteenth European Conference on Computer Systems*, pages 1–16, 2020.
50. Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3):50–60, 2020.
51. Anthony Liguori. The Nitro Project – Next generation AWS infrastructure. In *Hot Chips: A Symposium on High Performance Chips*, 2018.
52. Joshua Lind, Christian Priebe, Divya Muthukumaran, Dan O'Keeffe, Pierre-Louis Aublin, Florian Kelbert, Tobias Reiher, David Goltzsche, David M. Eyers, Rüdiger Kapitza, Christof Fetzer, and Peter R. Pietzuch. Glamdring: Automatic application partitioning for Intel SGX. In *Proceedings of the USENIX Annual Technical Conference*, pages 285–298, 2017.
53. Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. Recalibrating global data center energy-use estimates. *Science*, 367(6481):984–986, 2020.
54. Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R Savagaonkar. Innovative instructions and software model for isolated execution. *Hasp@ isca*, 10(1), 2013.
55. Rich Miller. The bitcoin energy debate: Lessons from the data center industry. https://datacenterfrontier.com/the-bitcoin-energy-debate-lessons-from-the-data-center-industry, 2021.
56. F. F. Moghaddam, M. Cheriet, and K. K. Nguyen. Low Carbon Virtual Private Clouds. In *Proceedings of the 2011 IEEE International Conference on Cloud Computing*, pages 259–266, Washington, D.C., USA, 2011.
57. Michael Naehrig, Kristin Lauter, and Vinod Vaikuntanathan. Can homomorphic encryption be practical? In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, pages 113–124, 2011.
58. United Nations. Sustainability. Website, November 2022.
59. Nate Nelson. Threat post: Bug in the linux kernel allows privilege escalation, container escape. https://threatpost.com/bug-linux-kernel-privilege-escalation-container-escape/178808/, 2023.
60. Intergovernmental Panel on Climate Change (IPCC). Climate change 2022: Mitigation of climate change. Technical report, April 2022.
61. Daniel Perez and Benjamin Livshits. Smart contract vulnerabilities: Vulnerable does not imply exploited. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1325–1341, 2021.
62. Marc Pilkington. Blockchain technology: principles and applications. In *Research handbook on digital transformations*, pages 225–253. Edward Elgar Publishing, 2016.
63. Energy Innovation Policy and LLC Technology. How much energy do data centers really use? https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/, 2020.
64. Wendover Productions. The carbon offset problem. https://www.youtube.com/watch?v=AW3gaelBypY, June 2022.
65. Mark Randolph and William Diehl. Power side-channel attack analysis: A review of 20 years of study for the layman. *Cryptography*, 4(2):15, 2020.
66. Akshat Rathi, Natasha White, and Demetrios Pogkas. Junk carbon offsets are what make these big companies 'carbon neutral'. https://www.bloomberg.com/graphics/2022-carbon-offsets-renewable-energy/, November 2022.
67. Chuangang Ren, Di Wang, Bhuvan Urgaonkar, and Anand Sivasubramaniam. Carbon-Aware Energy Capacity Planning for Datacenters. In *Proceedings of the 20th IEEE International Symposium on Modeling,*

Analysis and Simulation of Computer and Telecommunication Systems, MASCOTS '12, pages 391–400, Arlington, VA, USA, 2012.

68. Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, pages 199—212, New York, NY, USA, 2009. Association for Computing Machinery.
69. Sam Steers. How crypto mining affects data centre sustainability. https://datacentremagazine.com/data-centres/how-crypto-mining-affects-data-centre-sustainability, 2021.
70. Sam Steers and Harry Menear. Carbon neutrality is a myth in the data centre industry. https://datacentremagazine.com/data-centres/carbon-neutrality-is-a-myth-in-the-data-centre-industry, April 2022.
71. The White House. Federal Sustainability Plan: Catalyzing America's Clean Energy Industries and Jobs. https://www.sustainability.gov/pdfs/federal-sustainability-plan.pdf, 2021.
72. Petar Tsankov, Andrei Dan, Dana Drachsler-Cohen, Arthur Gervais, Florian Buenzli, and Martin Vechev. Securify: Practical security analysis of smart contracts. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 67–82, 2018.
73. U.S. Environmental Protection Agency. Definition | CO2e. https://www3.epa.gov/carbon-footprint-calculator/tool/definitions/co2e.html.
74. Zhenghong Wang and Ruby B Lee. New cache designs for thwarting software cache-based side channel attacks. In *Proceedings of the 34th annual international symposium on Computer architecture*, pages 494–505, 2007.
75. Yuanzhong Xu, Weidong Cui, and Marcus Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *2015 IEEE Symposium on Security and Privacy*, pages 640–656. IEEE, 2015.
76. Jie You, Jae-Won Chung, and Mosharaf Chowdhury. Zeus: Understanding and optimizing GPU energy consumption of DNN training. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 119–139, Boston, MA, April 2023. USENIX Association.
77. Rui Zhang, Rui Xue, and Ling Liu. Security and privacy on blockchain. *ACM Computing Surveys (CSUR)*, 52(3):1–34, 2019.
78. Jianping Zhu, Rui Hou, XiaoFeng Wang, Wenhao Wang, Jiangfeng Cao, Boyan Zhao, Zhongpu Wang, Yuhui Zhang, Jiameng Ying, Lixin Zhang, and Dan Meng. Enabling rack-scale confidential computing using heterogeneous trusted execution environment. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1450–1465, 2020.
79. Dimitrios Zissis and Dimitrios Lekkas. Addressing cloud computing security issues. *Future Generation computer systems*, 28(3):583–592, 2012.

S. R. Hussain is an Assistant Professor in the Computer Science and Engineering Department at Pennsylvania State University. His research interests include systems and network security, formal methods, and sustainability. Hussain received the Ph.D. degree in Computer Science from Purdue University, West Lafayette. He is a Member of ACM and IEEE. Contact him at hussain1@psu.edu.

P. McDaniel is the Tsun-Ming Shih Professor of Computer Sciences in the School of Computer, Data & Information Sciences at the University of Wisconsin-Madison. McDaniel's research focuses on a wide range of topics in computer and network security and technical public policy, with interests in mobile device security, the security of machine learning, systems, program analysis for security, sustainability and election systems. McDaniel is a Fellow of IEEE, ACM and AAAS, a recipient of the SIGOPS Hall of Fame Award and SIGSAC Outstanding Innovation Award, and the director of the NSF Frontier Center for Trustworthy Machine Learning. He also served as the program manager and lead scientist for the Army Research Laboratory's Cyber-Security Collaborative Research Alliance from 2013 to 2018. Prior to joining Wisconsin in 2022, he was the William L. Weiss Professor of Information and Communications Technology and Director of the Institute for Networking and Security Research at Pennsylvania State University. Contact him at mcdaniel@cs.wisc.edu.

A. Gandhi is an Associate Professor in the Computer Science Department at Stony Brook University. His research interests include performance analysis, modeling, and evaluation; distributed systems; and sustainability. Gandhi received the Ph.D. degree in Computer Science from Carnegie Mellon University. He is a Senior Member of ACM and a Senior Member of IEEE. Contact him at anshul@cs.stonybrook.edu.

K. Ghose is a SUNY Distinguished Professor of Computer Science at SUNY Binghamton (Binghamton University). His research interests include energy-aware systems at all scales, processor microarchitectures and hardware security. Ghose received the Ph.D. degree in Computer Science from Iowa State University. He is a Member of ACM and IEEE. Contact him at ghose@binghamton.edu.

K. Gopalan is a Professor in the Computer Science Department at Binghamton University. His research interests are in computer systems including virtualization, security, operating systems, networks, and sustainability. He received his Ph.D. degree in Computer Science from Stony Brook University. He is a senior member of IEEE and a member of ACM. Contact him at kartik@binghamton.edu.

D. Lee is an Assistant Professor in the Computer Science Department at Stony Brook University. His research interests include compilers, operating systems, computer architecture, security, and sustainability. Lee received the Ph.D. degree in Computer Science Engineering from Michigan University, Ann Arbor. He is a Member of ACM and IEEE. Contact him at dongyoon@cs.stonybrook.edu.

Y. D. Liu is a Professor in the Computer Science Department at Binghamton University. His research interests include programming languages; software engineering; formal methods for security; sustainable and energy-aware applications and systems. Liu received the Ph.D. degree in Computer Science from the Johns Hopkins University. He is a Member of ACM. Contact him at davidl@binghamton.edu.

Z. Liu is an Associate Professor in the Department of Applied Mathematics and Statistics and Department of Computer Science at Stony Brook University. His research interests include optimization; machine learning; big data systems; sustainable and energy-aware applications and systems. Liu received the Ph.D. degree in Computer Science from California Institute of Technology. He is a Member of ACM. Contact him at zhenhua.liu@stonybrook.edu.

S. Mu is an Assistant Professor in the Department of Computer Science at Stony Brook University. His research interests include distributed systems, multicore systems, and sustainable systems. Mu received the Ph.D. degree in Computer Science from Tsinghua University. He is a Member of ACM. Contact him at shuai@cs.stonybrook.edu.

E. Zadok is a Professor at Stony Brook University. His research interests include computer systems, storage systems, security, performance optimizations, and sustainability. Zadok received the Ph.D. degree in Computer Science from Columbia University. He is a Senior Member of the IEEE Computer Society, an ACM Distinguished Member, and a member of USENIX. Contact him at Erez.Zadok@stonybrook.edu.