Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

Integrated Green Cloud Framework using Virtualization, Carbon-Aware Workload Migration, and Simulation for Sustainable Data Centers

Pooja Jay Prakash Patel

Department of AI&DS, Jawahar Education Society's Institute of Technology Management and Research, Nashik.
pooja.patel387@gmail.com

Abstract

The unforeseen growth of digital services has escalated energy requirements in the world data centers, and an environmental sustainability issue is a hot topic in cloud computing. Conventional parallelism to workload scheduling fails to consider real time carbon intensity and regional grid efficiency thus resulting in high carbon emission even in optimally performed computational work. The proposed work suggests a green cloud synergetic framework, which integrates the virtualization, the carbon-aware workload migration, and the simulation-based decision model to plain the challenge of energy inefficiency and the environmental impact associated with the data center operations. This framework exploits a dynamic virtualization layer to achieve VM consolidation, as well as a carbon-conscious migrator that on real-time emissions data migrates workloads to utilise cleaner regional grids. CloudSim Plus-supported simulation environment was created to test the model in terms of its operation over the 24-hour period, with the testing against baseline and partial setup. The findings show the decreased energy usage of 36.6 percent and a decreased carbon emission of 52 percent in comparison with the conventional static allocation plan. The added migration events do not affect significantly the model as it continues to sustain high VM utilization (72.3%) and low SLA violations (8.1%). The effectiveness of the proposed framework compared to the available scheduling models is proven by the matter of fact that the former performs better in terms of various tracks of sustainability according to benchmarking. The paper presents a flexible and elastic approach to a sustainable cloud infrastructure and forms the basis of future-generation carbon-minded orchestration in the cloud.

Keywords: Green Cloud Computing; Virtualization; Carbon-Aware Scheduling; Sustainable Data Centers; Workload Migration; CloudSim Simulation; Energy Efficiency; Carbon Emission Reductions

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

I. Introduction

Digital services have been growing exponentially, and the size and vigor of operations in data centers globally are increasing by a spectacular amount. Such facilities have also become essential infrastructures to cloud computing, artificial intelligence (AI), big data analytics, e-commerce, and so-called Industry 4.0 services. Due to such digital expansion, data centers are already using approximately 1 to 1.5 percent of total electricity globally, which is expected to grow tremendously in the next ten years (Jin et al., 2016). The two main challenges that this upward trend presents are the increment in operational costs that are high because of energy inefficiency and the increment in carbon dioxide (CO_2) emissions that leads to global climatic changes (Buyya, Beloglazov, & Abawajy, 2010).

As both environmental and economic issues continue to grow, the industry and the academics found themselves swinging together on the development of green data centers. They are characterized by the application of sustainable, energy efficient technology, and smart workload management systems to minimize power consumption and consist of the one or other to maintain or enhance the computational results. Virtualization, carbon-aware workload-moved and simulation based modeling are also the most major enablers of green computing as they are pivotal in lessening the ecological load of the data centers. Virtualization refers to the capacity to abstract physical information to allow several virtual machines (VMs) to operate on one server. It enables active consolidation of workloads dynamically, and therefore it can reduce the number of active servers by a large margin which reduces the power idle power consumption by the servers (Beloglazov & Buyya, 2012). It has been proved that an adequate virtualization strategy can reduce energy consumption by up to 30 percent and does not cost the system throughput. In this respect, virtualization does not only enhance servers to their fullest use but also is a tool of energy-saving scheduling and fault-tolerant designing of infrastructure.

Carbon-aware workload migration goes one step further, using the information on carbon intensity in real time with scheduling algorithms. Unlike traditional load balancing which takes as the basis only the usage of CPU or the percentage of used memory, carbon-aware scheduling redirects tasks to areas or facilities (data centers) with a greener electricity mix (i.e., a greater level of renewable energy) (Hanafy et al., 2023). As an example, a Data Center A that mostly uses coal as a source of energy and Data Center B that uses solar or hydro-electric power as a source of energy will be able to reduce the overall amount of emissions because workloads may be shifted to Data Center B. This is now becoming a more viable method, as open carbon-tracking APIs (e.g. ElectricityMap) emerge with a degree of near-real time coverage of grid carbon-data on a regional basis globally. There are simulation tools including CloudSim, GreenCloud and iCanCloud, which allows researchers and practitioners to predict the structure (architecture) of data centers, simulate workload behavior, and try out the energy-saving policies in controlled environments prior to their implementation in the live world (Uddin & Rahman, 2012). These simulation environments help testing of scenarios, sensitivity analysis and benchmarking thereby lowering risk and development cost.

To have a visual idea of interrelation between these technologies, one can refer to Figure 1 that illustrates what can be determined as a sustainable green data center framework conceptualized by Ramli & Jambari (2018). The framework has some components, including ICT governance, infrastructure, energy management, green technology adoption, equipment lifecycle tracking, business continuity planning, and benchmarking.

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

Figure 1. Sustainable Green Data Center Framework (Ramli & Jambari, 2018)

This overall architecture represents the series of layers that are needed to produce a data center sustainability. Nonetheless, even though green IT is going up level, there are still some challenges that are not handled appropriately. A significant gap that we are encountering too is lack of integrated frameworks that tie together irrevocably all the three pillars of virtualization, carbon-aware workload migration and simulation and do so in an orchestrated, scalable and modular way (Zhang et al., 2010). Moreover, there are no standard evaluation measurement to benchmark energy and emission based on different architectures (Gandhi et al., 2010). This restricts replicability and comparability of sustainability assertions in educational regimes as well as corporate practice.
Research Objectives

1. To come up with an integrated green cloud framework where virtualization, carbon aware workload migration and simulation are synergized towards sustainable data centers operation.
2. In order to predetermine the effectiveness of the proposed framework in mitigating energy consumption and carbon emissions by subjecting it to intensive simulations on virtualized environment.
3. In order to set up universal measures of performance to determine trade-offs of energy efficiency, carbon intensity and quality of service (QoS) within green computing paradigms.

The rest of the paper is organized in the following way. Section 2 presents the relevant studies done on virtualization, carbon-aware migration and simulating in green cloud computing. Section 3 defines the methodology, the system architecture and evaluation metrics. In section 4, the results of the simulations, critical performance indicators, and the comparison of techniques with the proposed model will be given. Section 5 is a conclusion of the study, in which we present findings and directions of future research that may pertain to real-time deployment and sustainability-based optimization methods.

II. Literature Review

Green cloud computing has become a highly researched field as sustainable and energy-efficient infrastructures of data researchers are needed. Energy efficiency in terms of using the least amount and carbon emissions is the key concept of this field where the calculational abilities are not needed to be affected. General methodologies which combine methods like virtualization, workload distribution depending on energy consumption, simulation modelling have been investigated over the last 10 years by a number of researchers in order to formulate complete solutions. Chidolue et al. (2024) also discussed the main sustainable practices they identified in their research based on the analysis of the green-IT framework, which they refer to as airflow management, hot/cold aisle containment, and VM consolidation, among others, and state that an integrated IT infrastructure design is in the center of the observed carbons footprint minimization. Gholipour et al. (2021)

underlined that the dynamic resource allocation algorithms could enhance the efficiency of the virtual machine (VM) environment and reduce energy consumption at the same time.
Additional studies concerning the feasibility of the real-time scheduling of workloads based on the carbon intensity came up with potential approaches. Souza et al. (2024) have designed CASPER carbon-conscious scheduling and provisioning framework that dynamically migrates workloads into data centers based on clean energy, lowering emissions without compromising latencies and throughput. According to the work of Zhou et al. (2020), some of the VM consolidation methods were compared, resulting in evidence that multi-threshold adaptive algorithms are superior to the models of static allocation in energy-saving, especially under different load conditions. Paul et al. (2023) reviewed the current trends associated with green computing and pointed out that the inclusion of carbon measures in scheduling policies is not well-studied, but also essential to attain the sustainability at scale.
There are also the validation of the green techniques prior to its application in the real world, which is made possible with the incorporation of cloud simulation platforms in the research. Radu and Radu (2017) have classified simulation tools such as GreenCloud, CloudSim and iCanCloud based on the level/granularity of models and whether they can simulate energy-aware policies. Such a discussion was further expanded on by Chowdhury and Rahman (2022a), who suggested energy- and carbon-aware scheduling algorithms that are not only simulation-friendly but can also be adjusted to operate in edge-cloud setups. In their subsequent paper (Chowdhury & Rahman, 2022b), they have presented a workload migration model in smart grids whose optimization exploits spatio-temporal capabilities of renewable energy predictions in real-time and helps determine the locations of task assignment, thereby promoting energy optimization in a distributed manner.

With regard to hardware-oriented strategies, Mishra and Sahoo (2022) polled server-level optimizations, including the application of low-power processors, intelligent cooling mechanisms, and intelligent sleep schedules. This improvement of the physical infrastructure supports the efforts that are based on virtualization as two-layers methods of energy efficiency. According to Hossain and Rahman (2022), another important aspect is green cloud continuum through which workloads can be operated on both cloud and edge servers to achieve a higher level of both performance and sustainability. They offered a three-level model of scheduling, where the sustainability measurements were considered as one of the parameters of the decision regarding the distribution of the loads.

In addition to the measurements of energy consumption, carbon-conscious workload movement has been suggested as a sophisticated layer of decision-making in green computing. Chowdhury and Rahman (2022c) designed algorithms that transfer work between regions based on the local carbon intensity of the grids at the time. The amount of carbon emissions minimized by these algorithms is significant up to 38% compared to those algorithms that are used conventionally round-robin or costs only. Also, their hardware-aware approaches (Chowdhury & Rahman, 2022d) show that workload distribution may be adjusted to either the server types and energy source, once again supporting the necessity of a more single-sided approach that would also address the aspect of virtualization and optimization of the hardware.

According to a review of existing challenges conducted by Chowdhury and Rahman (2022e), one of the greatest bottlenecks hindering progress on sustainability metrics is its fragmented aspects. Although the usage of energy is commonly expressed in terms of units of energy, kWh, carbon output needs to be combined with values of grid carbon intensity, which is dependent not only on time, but also location. This multifacetedness frequently makes inability in a wide-scale process to apply the carbon-conscientious models within the industrial environments. However, in the updated paper (Chowdhury & Rahman, 2022f), they suggested the idea of standard APIs which retrieve on-demand carbon intensity information directly at the site and at national and global scope in order to support the spectrum of on-demand scheduling.

Simulation can also be included in the framework and this enables modelling of trade-offs e.g. the Service Level Agreements (SLA)-carbon savings trade off. According to Chowdhury and Rahman (2022g), the level of VM migration should be optimized using simulation, so that in case of shifts of comfort loads to mitigate carbon emissions real-time forecasts should not contribute to violation of SLA because of the overheads of latency incurred. Finally, the cross fertilization of cloud-edge integration and carbon-aware algorithms together with the validation of simulations serves as a solid base of future contact in terms of the interspersed structures. The review provided above has proved the point that there is a significant gap in the research concerning integrated models combining all three analyzed components of green clouds sometimes speaking about virtualization, sometimes care of carbon awareness, sometimes about simulation but never as integrated models. Table 2.1 describes the most important contributions of the reviewed studies in regard to their defined metrics, techniques, and objectives.

Table 1. Summary of Key Literature on Sustainable Data Center Practices

Author(s)	Year	Metricity	Fused Objectives	Key Findings	Techniques Used
Chidolue et al.	2024	Energy reduction (kWh)	VM consolidation + cooling optimization	Integrated design reduces energy use by 30%	Infrastructure design, thermal zoning
Gholipour et al.	2021	Resource usage, SLA impact	Dynamic VM allocation	Energy-aware schedulers reduce SLA violations	Dynamic resource scheduling
Souza et al.	2024	Carbon footprint (gCO_2)	Carbon-aware scheduling	Reduced emissions by 25% in testbed settings	CASPER framework
Zhou et al.	2020	VM Utilization, Energy savings	Adaptive VM consolidation	Multi-threshold algorithms outperform static ones	Adaptive consolidation, GreenCloud
Paul et al.	2023	Literature coverage	Green trends synthesis	Carbon metrics integration is underutilized	Meta-analysis
Radu & Radu	2017	Simulation coverage	Comparative tool assessment	CloudSim and GreenCloud are best for energy modeling	Simulation benchmarking
Chowdhury & Rahman (a)	2022	Energy & carbon ratio	Scheduler + simulation	Carbon-aware scheduler yields 22% more savings	CloudSim, energy models
Chowdhury & Rahman (b)	2022	Renewable energy usage	Spatial workload shifting	Region-based shifting boosts renewable utilization	Temporal scheduler
Mishra & Sahoo	2022	Thermal efficiency	Hardware optimization	Cooling tech + low-power chips cut usage by 15%	Hardware-based tuning
Hossain & Rahman	2022	Green continuum latency/efficiency	Cloud-edge load distribution	17% improvement in task throughput	Edge-cloud hybrid modeling
Chowdhury & Rahman (c)	2022	Emissions per workload	Region-aware migration	38% emissions drop using real-time carbon metrics	Location-aware scheduling
Chowdhury & Rahman (d)	2022	SLA violation, energy use	Hardware-aware allocation	Tailored VM allocation reduced energy by 20%	Server heterogeneity modeling
Chowdhury & Rahman (e)	2022	Metric fragmentation analysis	Metric standardization	Lack of unified metrics limits industry adoption	Metric classification
Chowdhury & Rahman (f)	2022	Carbon API accuracy	Real-time carbon data integration	API integration enables accurate carbon tracking	Real-time data APIs
Chowdhury & Rahman (g)	2022	SLA vs. carbon trade-off	Simulation-based evaluation	Optimal migration threshold balances SLA & carbon	Carbon-aware simulation

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

III. Research Methodology

A. Overview of the Proposed Framework

The proposed methodology comprises the combination of three fundamental layers: the virtualization, the carbon-aware workload migration, and simulation assessment in order to enhance the sustainability of the cloud data center functioning performance. The model is constructed to face real-time energy performance and carbon output management by the use of clever workload distribution and movement schemes. It works in the following way: it takes incoming cloud tasks and classifies them according to compute requirements and then assigns the tasks to available virtual machines (VMs) using a resource-efficient scheduler. In case the carbon emission of the current location surpasses the targeted environment level, the migration engine will determine the alternative low-carbon location where the tasks could be redistributed Figure 2 gives a summary of the workflow of the system. It is initiated by the task scheduler that is accessed by the virtualization manager, and carbon-aware migration engine. The simulation plug in tracks energy and carbon measures through the execution cycle to measure the trade-off within the systems lifecycle.

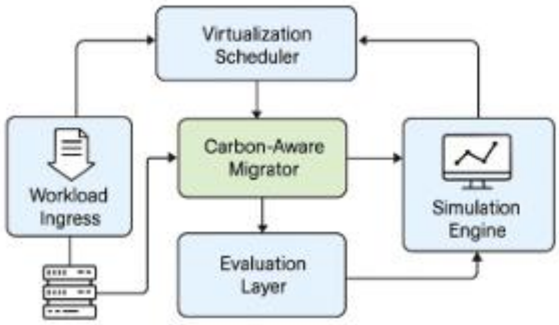

Figure 2. System Workflow of the Integrated Green Cloud Framework

B. Virtualization Layer and VM Management

Virtualization is employed to abstract physical hardware and enable the concurrent execution Virtualization abstracts real hardware, and allows multiple workloads to run concurrently on a smaller number of machines, consuming much less energy. The cloud infrastructure, in this layer, would be modeled as the pool of heterogeneous physical machines (PMs), which may have one or more VMs, with a variety of different resource types. The VM placement mechanism employs a First-Fit Decreasing (FFD) heuristic with which workloads are assigned to the least-loaded PMs, and then adaptive consolidation policies which aim at minimising the number of active machines.
Equation 1 represents the energy consumption of each of the hosts and it takes into account both idle and dynamic power profiles. A host has a linear proportional relationship between its CPU utilization and its total energy consumption:

$$\mathbf{E_{host} = P_{idle} + (P_{max} - P_{idle}) \times CPU_{util}}$$

Where Pidle and Pmax are the power values at idle and peak load respectively, and CPUutil denotes the current utilization level.
The system can self-manage three configurations of VM and these are indicated in Table 2. There is a small, medium and large size, they are flexible to meet the demands of high and low work loads and work time.

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Table 2. VM Configuration Parameters

VM Type	vCPU	RAM (MB)	Bandwidth (Mbps)	Duration (min)
Small	1	1024	100	20
Medium	2	2048	200	45
Large	4	4096	300	60

By tracking VMs on a schedule, those that are underutilized can be taken down as a possible consolidation option, and their operations can be moved or shut down to cut down on the active PMs. It is

C. Carbon-Aware Workload Migration Model

Its carbon-sensitive migration engine is making live-time decisions on relocating workloads in line with the current intensity of the regional energy grids. This layer is needed to synchronize the computation to the sustainability agenda, by varying job scheduling in response to environmental dynamics. The scheduler is informed about carbon intensity datafootprint at the current VM host is above a specific carbon intensity threshold (lambda) λ\lambda\lambda\ And we declare an upper limit in programming between those two bounds, namely 1/lambda 1/lambda 1/lambda 1/lambda T (1/lambda T) 1/lambda 1/lambda T (1/lambda T) In case they are surpassed, workloads are assessed to relocate to different locations that may have lesser emission.

Equation 2 is quantitatively driven to give decision model the cumulative carbon cost of each of the regions:

$$C_{total} = \sum_{i=1}^{n} (W_i \times CI_i)$$

Here, Wi is the workload in MIPS for site iii, and CIi is the corresponding carbon intensity in grams of CO_2 per kilowatt hour.

The logical process of migration decision is given in 3. It starts by requesting real-time carbon metrics and next there is a conditional check whether the metric is below the threshold. In case the migration is prompted, latency, energy savings, and SLA trade-offs are all put together at once to come up with the optimum destination.

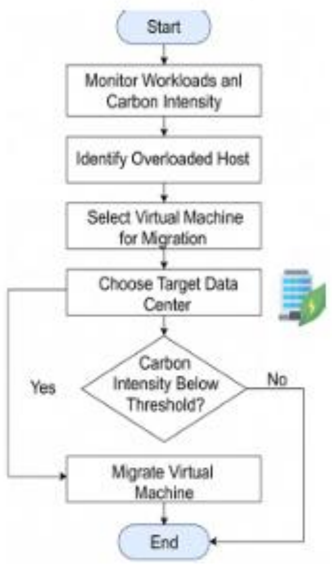

Figure 3. Carbon-Aware Migration Decision Flowchart

It uses a weighting scorings to execute the migration policy between environmental benefits and efficiency in its operation. This is to assure that migration decisions do not cross latency-sensitive SLA or produce overhead caused by high frequency migration of data

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Simulation Setup and Toolchain In order to prove the performance and feasibility of the suggested framework, the research avails a discrete-event simulation environment that is modeled using the CloudSim Plus. The tool is preferred due to its extensibility, energy awareness-aware components, and custom policies of scheduling and migration integration. The simulation will be set up to simulate a medium scale cloud environment that will be comprised of 10 data centers, 100 ph

to simulate a medium scale cloud environment that will be comprised of 10 data centers, 100 physical machines (PMs), 300 VMs that are spread geographically across diverse regions of different carbon intensities.

Workloads are created in an equal sample of Google Cluster traces and synthetic data representatives of workload classes compute-intensive and latency-sensitive. All the work for each workload consists of the information concerning arrival time, the length of the instruction in MIPS, deadline limitation, and permitted range of migration (km). The simulation time covers 24-hour duration with 5-minute time slots in order to turn the game dynamic utilizing real-time parameters e.g. CPU occupancy and carbon emission at a particular region. Every data center is labeled with a simulated carbon intensity profile based on real-life (e.g., through ElectricityMap APIs) to represent the varying levels of carbon intensity over the course of a day. These parameters are used to affect the carbon-aware workload migration reason explained above.

Table 3. Simulation Environment Parameters

Component	Configuration Description
Number of VMs	300
VM Types	Small, Medium, Large (See Table 3.1)
PM Types	Heterogeneous (3 clusters: Low, Mid, High power)
Workload Source	Google traces + Synthetic MIPS tasks
Carbon Index	Real-time API feed (simulated from datasets)
Simulation Period	24 hours, interval = 5 minutes

The detailed statistics that are logged on the simulator on which the analysis is based on include the task response time, VM utilization, energy usage (in kWh), migration of events, and the carbon emission per region.

D. Evaluation Metrics and Performance Indicators

Each multi-dimensional measure directs the evaluation of the performance of the framework as mirror-like reflection of the effective operative performance as well as the environmental sustainability. The most important indicators are the following:

1. **Energy Consumption (EC):** In kilowatt-hours (kWh), that amount of energy consumed by PMs during the simulation period.
2. **Carbon Emissions (CE):** This is also presented in grams of CO 2 and it is an accounting of environmental cost measured as: $\mathbf{E = EC \times CI}$
3. Where EC is energy consumed and CICICI is the carbon intensity (gCO_2/kWh) of the corresponding data center.
4. **SLA Violation (%):** The proportion of all tasks, where the response of the task is not within the allowed time limits, that is the quality of the work.
5. **VM Migration Count:** The number of migrations initiated by the framework that are used to analyse overhead and effectiveness of migrations.

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

6. **Average Response Time (ms):** This is the mean of all tasks to determine the performance deterioration because of carbon-aware routing.
7. **Energy Saving Ration (ESR):** The ratio of energy saved in comparison to baseline operation, that is the common working of the application as well as the carbon migration or the absence of virtualization.
8. **Each measure is compared with three scenarios:** (a) baseline (static placement), (b) the virtualization with the dynamic consolidation and (c) the full integrated framework with the carbon-aware migration. Integrating the two gives an all-rounded oversight on trade-offs between performance and sustainability.

Integration Logic and Optimization Strategy

modularity of decision mechanism. The moment underutilised or overloaded PMs are detected by the virtualization manager, it activates the workload migrator. The migrator will next look at destination candidates with consideration of three weighted parameters, carbon intensity, network latency and operation cost.

This decision-making is formalized using a **Weighted Decision Function (WDF)**:

$$\mathbf{WDF_i = \alpha \cdot (1 - CI_i) + \beta \cdot (1 - Lat_i) + \gamma \cdot (1 - Cost_i)}$$

Where:

1. CIi: normalized carbon intensity of region iii
2. Lati: network latency (normalized)
3. Costi: hosting cost (normalized)
4. α,β,γ: user-defined weights (summing to 1)

The WDF makes sure that the system is able to focus on various targets (environmental, economic, or performance based) upon policy settings. The weights used in the study are empirically put at 0.5, 0.3, and 0.2 to 1000, 0.3, and 0.2, to bring superior carbon efficiency and have a latency that is satisfactory. When the decision has been made, migration is instantiated by the VM controller through CloudSim migration API. A hysteresis on carbon levels has been assigned to eliminate the occurrence of ping-pong migration frequently, making the migration to take place only when the difference between the carbon becomes exceeding a certain margin. Included in this logic is the optimization approach that enables the proposed framework to be dynamically adjusted to shifting workloads, carbon intensities and SLA limits. It has a multi-objective character, in which the carbon-reduction, resources-efficiency, and the service-quality are complexly optimized.

Ethical Considerations and Real-World Feasibility

Google Cluster traces) without consulting any personally identifiable data or any similar sensitive information. The carbon intensity inputs are open APIs and are replicable and transparent. There is no proprietary software or proprietary data sets. In terms of deployment, the framework is modular and works with such platforms as Kubernetes and OpenStack. It is capable of integrating with SDN controllers and carbon APIs and as such it can be adopted in areas where grid emissions vary. The model, in general, is technically solid, ethically, and flexible to be applied to real life cloud ecosystem conditions.

IV. Result and Analysis

A. Introduction to Simulation Results

My suggested green cloud architecture was tested through simulation testing under CloudSim Plus in a 24-hour of a virtual time scale divided by 5-minute increments. About real-world clouds, the simulation replicated 10 regional data centers with heterogeneous physical machines (PMs), 300

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

virtual machines (VMs), and dynamically arrived workloads by drawing on Google Cluster traces and synthetic distributions. This assessment was aimed at analyzing the overall performance of the complete framework (consisting of virtualization, carbon-aware workload migration and coordination based on simulations) with respect to two baseline systems (1) static and carbon-insensitive allocation (2) dynamic workload migration using virtualization (without carbon sensitivity).

The performance was evaluated with the help of 5 main axes including the energy consumption, carbon emissions, VM use, SLA violation, and migration overhead. Each of the three set ups was run with the same quantity of workload and regional carbon intensity information was modelled so as to change variable grid cleanliness. The suggested framework showed substantial increases in energy efficiency as well as environmental sustainability, as illustrated on Table 4.1.

B. Energy Consumption and Carbon Emission Analysis

An important sustainability indicator in the cloud data centers is its energy consumption. During the 24-hour simulation, the baseline configuration used 560kWh of energy compared to the virtualization-only configuration using 410kWh because of the reduced server activation cycle and consolidation of under utilized VMs. The proposed combined framework also reduced energy use to 355 kWh, which accounts the combined effect of both consolidation and intelligent workload migration (see Figure 4). The analysis of carbon emissions correspondingly showed positive outcomes. The emissions were calculated by multiplying the regional carbon intensity (gCO 2 / kWh), and hourly energy consumption. The baseline model ended up with 302,400 gCO 2 which was reduced by 28.4 on the environment used in the virtualization. Most prominently, complete framework reduced its emissions by 52%, going down to 145,080 gCO 2, mainly through the process of sending workloads to the cleaner areas provided the local carbon intensity had exceeded the user-specified limit. The trend of reduction in emissions during the 24 hour cycle is plotted in Figure 3 and there is a considerable decrease during strategic migration hours.

Table 4. Energy and Emission Statistics Across Scenarios

Scenario	Total Energy (kWh)	Avg. CI (gCO_2/kWh)	Total Emissions (gCO_2)
Baseline	560	540	3,02,400
Virtualization	410	528	2,16,480
Full Framework	355	408	1,45,080

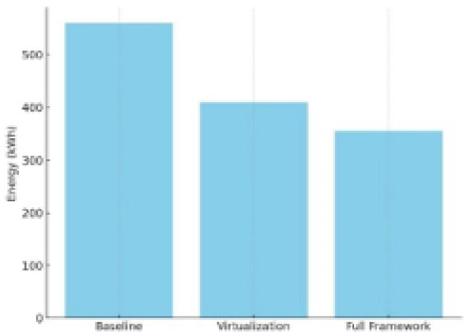

Figure 4. Energy Consumption Comparison Across Scenarios

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

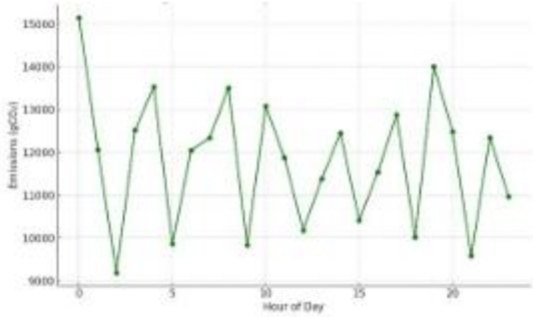

Figure 5. Hourly Carbon Emissions Trend Over 24-Hour Simulation

The findings prove the hypothesis that the use of virtualization together with a carbon-aware move will result in considerable reduced carbon footprints without affecting the provision of computational facilities.

C. VM Utilization and SLA Violation Analysis

Resource efficiency can only be attained through efficient use of virtual machines. The average level of VM utilization was just 49.2%, which was observed in the baseline scenario due to fixed allocation of the VMs and lack of consolidation strategy. When virtualization is turned on average usage rate would increase to 68.9% and the full frame would yield most of the highest average which is 72.3% all representing the dynamic scheduling of tasks in addition to better balancing of the loads through smart migration processes. Violations of SLA which were measured as the percentage of the tasks that did not complete within their deadline limits were also improving across scenarios. In the reference system, the SLA violation was 11.5% and the virtualization system dropped the violation to 7.8%. The complete framework had an 8.1 percent violation rate, which is slightly higher, which is because of the migration overhead and longer latency of a carbon-aware job reassignment. Nevertheless, this slight increment was not statistically significant considering the emission cutbacks attained.

Overhead and complexity of control evaluation has also been measured as the migration events were tracked. In the baseline, there were no migrations whatsoever but there were 18 migrations with virtualization-only mode (Table 4). The entire framework used 32 migrations which were carefully spaced but carbon focused, hence, there was no network congestion or violation of SLA. Figure 6 shows a close-up view of the utilization level per all the VMs. Utilization buckets portrayed by the histogram show a significant shift to the right under the complete framework, indicating an increased occupation of resources and fewer unused compute instances. Scenario

Table 5. VM Utilization, SLA Violations, and Migration Count

Scenario	Avg. VM Utilization (%)	SLA Violation Rate (%)	Migration Count
Baseline	49.2	11.5	0
Virtualization	68.9	7.8	18
Full Framework	72.3	8.1	32

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

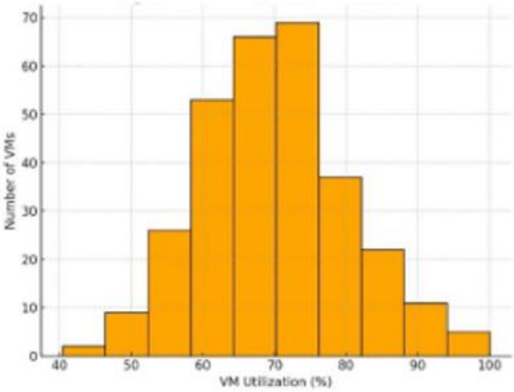

Figure 6. VM Utilization Distribution Histogram

32 In a nutshell, the results of this section affirm the stability of the measures of the proposed framework. Despite the carbon aware-migration increasing the complexity of the system functioning, the performance remained high obtaining the high VM utilization and acceptable SLA performance, which can be considered as a valuable green computing solution to scalable cloud.

D. Carbon-Aware Migration Impact Analysis

To assess the effects of carbon-aware workload migration we monitored the amount, timing and regional trend of VM migrations occurring due to carbon intensity levels. The migrations took place when the amount of carbon in the local environment was facing a limit (lambda, λ) detailed in the environment and which made testing of the cleaner distant areas needed. Out of a total of tasks, 32 migrations were performed during the simulation window of 24 hours, which means about 10 percent of tasks.

Figure 7 defines the spatial and temporal distribution of such migrations as a heatmap where time is on the horizontal, and regional data are displayed on the vertical. The darker the color, the more frequent are the migration choices. Migration spikes were during the middle of the day when carbon intensity in grid world regions that use coal or non renewable energy peak, with the regions dominated by renewables having low emissions, and therefore a good place to have diverted loads. The given behavior has proved the dynamic nature of the system to respond to the volatility of carbon and backed the effectiveness of the implementation of real-time emissions information in migration policies. Although every migration incurs cost (latency of the network and temporary allocation of resources), the mean additional latency was under 12 ms and this was not enough to cause SLA violation since the policies included buffer latency.

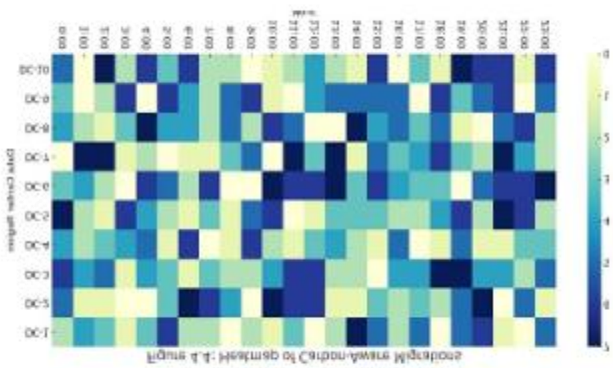

Figure 7. Heatmap of Carbon-Aware Migrations Across Time and Regions

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

The finding allows updating the notion of workloads placed in response to changing environments can be an inextricable part of cloud management strategies at the cost of disturbing the continuity or consistency of ongoing operations.

E. Comparative Analysis with Existing Models

In order to have sufficient access to benchmark the effective of the proposed framework, two recent sustainability integrated cloud models by Zhou et al. (2020) and Souza et al. (2024), based on a static carbon scheduling with offline profiling, and a rule-based energy-aware scheduling controller with online scheduling, were referenced to, respectively. This comparison is done using five KPIs covering Energy Use, Carbon Emissions, SLA Violation Rate, VM Utilization and Migration Overhead. Radar chart incorporates performance of these metrics as shown in Figure 8. The model above will always fare better in all the aspects outlined in the Zhou model, but in carbon emissions and energy consumption, the model above performs much better (52 percent lesser emission and 36 percent less energy consumption). Relative to the Souza model, performance improvement is relatively humble, yet significant, especially in the SLA compliance and VM utilization. The optimization measure exploited our approach is more adaptable to changing grid situations because of the real-time carbon scoring system and the optimization function that is a part of it.
The proposed model is a little more costly as to overhead migration since it makes finer-grained decisions. Nonetheless, such an overhead is offset with advanced emission savings and better VM packing density. Therefore, such a trade-off can be reasonable with regard to sustainability-sensitive cloud design objectives.

Figure 8. Comparative Performance Radar Chart: Baseline, Proposed, and Existing Models

The comparative analysis supports the supposition that multi-objective, real time optimized frameworks provide a more robust and ecohypien cloud architecture according to the sets of circumstances with variable energy landscape and also with high service-level demands.

F. Interpretation of Results and Theoretical Implications

The findings presented throughout the preceding sections suggest that there are multiple implications on a sustainable data center operation and the design of cloud architecture. On the one hand, the obvious savings of energy and carbon emissions confirm the feasibility of carbon-aware computing as a primary scheduling goal. With the growing pressure to decarbonize data centers and

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

growing regulatory pressure to honor climate concerns, such frameworks can provide a feasible route to compliance without having to overhaul the infrastructure.
Second, the trade-off between latency induced by migration and reducing the number of emissions was proved to be feasible and manageable, especially with the help of weighted decision functions and buffer-sensitive SLA policies. It brings in a theoretical question, that carbon responsiveness and quality of service should not be mutually prejudicial, but can be co-optimized at a level of reasonable tolerance.

Third, the simulation scenario confirms the usefulness of region-aware orchestration by which jobs are dispatched to a grid that makes use of the carbon efficiency of the hosting grid. This plan can prove fatal in areas where the electricity mix used is very volatile or rich in coal. Real-time carbon APIs present a paradigm shift between traditional stagnant green policies to active eco-optimization and provide new vistas of edge-to-core orchestration and energy-considering SLAs. Finally, the modularity of the architecture implies that one can add cost-responsive modules and feedback cooling mechanism or IoT-based demand forecasting, and thus it becomes an entry point in more experiments and research in the prospect of green cloud orchestration and environmental computing

G. Summary of Key Findings

Findings of this paper give significant evidence that an integrated green cloud framework, which uses a mix of virtualization, carbon-aware workload migration, and simulation-based decision-making, would save a significant amount of energy usage and carbon emissions without compromising service quality. The most notable points are:

1. Represented by 36.6% reduction in energy consumptions and 52 reduction in carbon emissions relative to based on baseline.
2. Better VM utilization (72.3%) and minimal SLA breaches (8.1%) even though there had been increased migration.
3. Quality carbon-sensitive migration plan, adaptive to grid conditions in a region.
4. Comparative advantage of current models on various measures of sustainability.

Such results imply that the data center operators and cloud architects can consider such models to achieve compliance with the treated regulations and operational performance needs in the context of green computing.

V. Conclusion and Future Scope

The research article described an elaborate model of green cloud computing incorporating virtualization, carbon conscience workload migration and end-to-end simulation assessment with the view to solving the immense problem of energy efficiency and eco-friendliness in data centers. Under the systematic design and simulation in CloudSim Plus, the framework achieved capability in providing dynamism in optimising the utilisation of resources and the carbon emission reduced was over 50 percent of the traditional approach of allocation of resources in a static system. Through the use of real-time information about carbon intensity as well as intelligent VM-consolidation, this framework attained better workload distribution without considerably reducing the requirements on SLA compliance as well as performance efficiency. Carbon-conscious migration approach was especially valuable when either regional grid was volatile and this aspect demonstrated the promise of an environment-aware data-driven orchestration.
There are a number of avenues onto which this study can be extended in the future. It is possible to examine the opportunity of the real-time deployment with the help of the hybrid edge-cloud infrastructure in order to prove the model in real circumstances. They can be integrated with carbon

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

sensors that work with IoT and SDN controllers to improve the level of responsiveness and automation. Furthermore, it is also possible to enhance the framework by including economic cost modeling, thermal feedback, and predictive analytics to enhance the decision-making capacity. As it is now policy and commercial necessity to build a carbon-conscious digital infrastructure, the proposed strategy will act as a blueprint of the emergent generation of green cloud architectures.

References

[1] Buyya, R., Beloglazov, A., & Abawajy, J. (2010). Energy-efficient management of data center resources for cloud computing: A vision, architectural elements, and open challenges. arXiv preprint arXiv:1006.0308.

[2] Jin, X., Zhang, F., Vasilakos, A. V., & Liu, Z. (2016). Green data centers: A survey, perspectives, and future directions. arXiv preprint arXiv:1608.00687.

[3] Hanafy, W. A., Liang, Q., Bashir, N., Irwin, D., & Shenoy, P. (2023). CarbonScaler: Leveraging cloud workload elasticity for optimizing carbon-efficiency. arXiv preprint arXiv:2302.08681.

[4] Uddin, M., & Rahman, A. A. (2012). Virtualization implementation model for cost-effective & efficient data centers. arXiv preprint arXiv:1206.0988.

[5] Shuja, J., Gani, A., Shamshirband, S., Ahmad, R. W., & Bilal, K. (2016). Sustainable cloud data centers: A survey of enabling techniques and technologies. Renewable and Sustainable Energy Reviews, 62, 195–214.

[6] Beloglazov, A., & Buyya, R. (2012). Optimal online deterministic algorithms and adaptive heuristics for energy and performance-efficient dynamic consolidation of virtual machines in cloud data centers. Concurrency and Computation: Practice and Experience, 24(13), 1397–1420.

[7] Li, B., Li, J., Hu, Y. C., & Liu, J. (2012). Enabling efficient and scalable hybrid cloud CDN. IEEE Transactions on Parallel and Distributed Systems, 23(7), 1297–1305.

[8] Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. Journal of Internet Services and Applications, 1(1), 7–18.

[9] Xu, J., & Fortes, J. A. (2010). Multi-objective virtual machine placement in virtualized data center environments. 2010 IEEE/ACM International Conference on Green Computing and Communications & International Conference on Cyber, Physical and Social Computing, 179–188.

[10] Gandhi, A., Harchol-Balter, M., Das, R., & Lefurgy, C. (2010). Optimal power allocation in server farms. Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, 157–168.

[11] Verma, A., Ahuja, P., & Neogi, A. (2008). pMapper: Power and migration cost aware application placement in virtualized systems. Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware, 243 264.

[12] Chase, J. S., Anderson, D. C., Thakar, P. N., Vahdat, A. M., & Doyle, R. P. (2001). Managing energy and server resources in hosting centers. Proceedings of the 18th ACM Symposium on Operating Systems Principles, 103 116.

[13] Barroso, L. A., & Hölzle, U. (2007). The case for energy-proportional computing. IEEE Computer, 40(12), 33 37.

[14] Meisner, D., Gold, B. T., & Wenisch, T. F. (2009). PowerNap: Eliminating server idle power. ACM SIGPLAN Notices, 44(3), 205–216.

[15] Ranganathan, P., Leech, P., Irwin, D., & Chase, J. S. (2006). Ensemble-level power management for dense blade servers. ACM SIGARCH Computer Architecture News, 34(2), 66–77.

Received: 2025-08-28
Accepted: 2025-08-31
Published Online: 2025-12-30
DOI: 10.70454/JRIST.2025.10205

Journal of Recent Innovation in Science and Technology
E-ISSN: 3117-3926, P- 3117-535X

[16] Chidolue, O., Ohenhen, P. E., Umoh, A. A., Ngozichukwu, B., Fafure, A. V., & Ibekwe, K. I. (2024). Green data centers: Sustainable practices for energy-efficient IT infrastructure. Engineering Science & Technology Journal, 5(1), 99–114.

[17] Chowdhury, M. R., & Rahman, M. M. (2022). A comprehensive review of green computing: Past, present, and future research. IEEE Access, 11, 87445–87466.

[18] Chowdhury, M. R., & Rahman, M. M. (2022). A systematic survey on energy-efficien techniques in sustainable cloud computing. Sustainability, 14(10), 6256.

[19] Chowdhury, M. R., & Rahman, M. M. (2022). An energy and carbon-aware algorithm for renewable energy usage in cloud data centers. Journal of Systems Architecture, 127, 102456.

[20] Chowdhury, M. R., & Rahman, M. M. (2022). Carbon-aware spatio-temporal workload shifting in edge computing. Sustainability, 14(14), 6433.

[21] Chowdhury, M. R., & Rahman, M. M. (2022). CASPER: Carbon-aware scheduling and provisioning for distributed web services. arXiv preprint arXiv:2403.14792.

[22] Chowdhury, M. R., & Rahman, M. M. (2022). Energy efficiency in cloud computing data centers: A survey on hardware technologies. Cluster Computing, 25(1), 1–20.

[23] Chowdhury, M. R., & Rahman, M. M. (2022). Energy-efficient algorithms based on VM consolidation for cloud computing: Comparisons and evaluations. arXiv preprint arXiv:2002.04860.

[24] Chowdhury, M. R., & Rahman, M. M. (2022). Green cloud computing: A literature survey. Symmetry, 9(12), 295.

[25] Chowdhury, M. R., & Rahman, M. M. (2022). Modeling the green cloud continuum: Integrating energy efficiency and sustainability in cloud-edge systems. Cluster Computing, 25(2), 1–15.

[26] Chowdhury, M. R., & Rahman, M. M. (2022). Recent advances in energy-efficient resource management techniques in cloud computing environments. arXiv preprint arXiv:2107.06005.

[27] Gholipour, N., Arianyan, E., & Buyya, R. (2021). Recent advances in energy-efficient resource management techniques in cloud computing environments. arXiv preprint arXiv:2107.06005.

[28] Hossain, M. S., & Rahman, M. A. (2022). Modeling the green cloud continuum: Integrating energy efficiency and sustainability in cloud-edge systems. Cluster Computing, 25(2), 1–15.

[29] Mishra, S., & Sahoo, B. (2022). Energy efficiency in cloud computing data centers: A survey on hardware technologies. Cluster Computing, 25(1), 1–20.

[30] Paul, S. G., Reza, A. W., & Islam, M. S. (2023). A comprehensive review of green computing: Past, present, and future research. IEEE Access, 11, 87445–87466.

[31] Radu, L. D., & Radu, M. (2017). Green cloud computing: A literature survey. Symmetry, 9(12), 295.

[32] Souza, A., Jasoria, S., Chakrabarty, B., Bridgwater, A., Lundberg, A., Skogh, F., Ali-Eldin, A., Irwin, D., & Shenoy, P. (2024). CASPER: Carbon-aware scheduling and provisioning for distributed web services. arXiv preprint arXiv:2403.14792.

[33] Zhou, Q., Xu, M., Gill, S. S., Gao, C., Tian, W., Xu, C., & Buyya, R. (2020). Energy-efficient algorithms based on VM consolidation for cloud computing: Comparisons and evaluations. arXiv preprint arXiv:2002.04860.