Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

Sustainable Cloud Computing Practices for Insurance Data Centers: Reducing Carbon Footprint While Scaling Digital Transformation

Sulabh Jain

Berkshire Hathaway Homestate Companies, USA

ARTICLE INFO

Received: 06 Oct 2025

Revised: 24 Nov 2025

Accepted: 06 Dec 2025

ABSTRACT

The trajectory of digital transformation in the insurance sector intersects critically with environmental sustainability imperatives as organizations navigate the transition away from legacy on-premises infrastructure toward cloud-based computing environments. Traditional data center operations are an environmentally burdensome resource utilization model due to inefficient resource use, high energy demands for computational needs and cooling, and significant water requirements, contributing disproportionately to organizational carbon footprints. Conversely, cloud computing architectures have the transformative property of reducing environmental impact, being hyperscale efficient with the capability of integrating renewable sources, and the ability to manage resources dynamically. Microservices decompositions, serverless and intelligent auto-scaling protocols can be used to achieve strategic migration policies that allow resource assignments that are far finer-grained to actual workload demands, and therefore are far less susceptible to the energy waste associated with fixed capacity infrastructure. Machine learning technologies and artificial intelligence techniques break down ongoing ethics of the environment through the rapidity of continuous optimization of environmental work by anticipatory workforce management, automated resource lives, and the creation of carbon-centric job variations matching computational tasks with periods of best grid carbon utility. Extensive measurement systems comprise Power Usage Effectiveness ratios, multiscope carbon accounting, and real-time monitoring dashboards, which show transparency, which is so important in making an informed decision and accountability. An association of sustainability measures with the corporate governance frameworks pertains to climate-responsible strategic technology determination as well as fulfilling the growing regulatory, shareholder, and customer requirements of climate responsibility.

Keywords: Cloud Computing Sustainability, Carbon Footprint Reduction, Green IT Infrastructure, Energy-Efficient Data Centers, Environmental Cloud Management

1. Introduction

The transition of the insurance industry to cloud computing has conventionally been considered from the perspectives of operational efficiency, scalability, and competitive advantage. However, environmental sustainability has emerged as a very critical dimension in this transformation. As the sector accelerates digital initiatives, the carbon footprint of IT infrastructure has turned out not only to be a regulatory concern but also a strategic imperative. Cloud adoption, if done thoughtfully, will provide an opportunity to bring energy consumption down dramatically while furthering business objectives. According to research, cloud computing can help decrease carbon emissions by approximately 88% compared with traditional on-premises data centers, while organizations can

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

realize energy efficiency improvements of 65% via optimized resource utilization and dynamic workload management.

Legacy on-premises data center infrastructure presents significant environmental concerns for the insurance industry. Conventional facilities require large amounts of electrical power for computation and cooling systems, and are often at less than optimal efficiency, with average PUEs greater than 2.0, reflecting that the same amount of energy is being used for cooling and auxiliary systems as by the computing equipment itself [1]. Utilization rates of servers in traditional enterprise data centers are often less than 20%, meaning computational resources are idle about 80% of the time but continue to draw baseline electrical power. Fixed infrastructure can be inflexible; situations arise when capacity is provisioned to meet peak demand but otherwise remains underutilized, which is a persistent waste of energy. Insurance carriers maintaining distributed data center networks face compounded inefficiencies, as smaller facilities cannot achieve economies of scale available to hyperscale operations [2].

Cloud computing architectures present a fundamentally different environmental equation. Properly executed cloud migration strategies achieve carbon footprint reductions ranging from 60% to 80% compared to equivalent on-premises infrastructure, with some implementations demonstrating reductions up to 98% when leveraging renewable energy sources [1]. These gains derive from multiple converging factors: hyperscale data center facilities operate at significantly higher efficiency ratios, leveraging advanced cooling technologies, optimized power distribution systems, and high-density server configurations that maximize computational output per unit of energy consumed. Modern cloud data centers achieve PUE values as low as 1.1, representing a 45% improvement in energy efficiency compared to traditional facilities [2]. Geographic diversity of cloud regions enables workload distribution strategies that prioritize locations with cleaner energy grids or favorable climate conditions, reducing cooling requirements. Beyond infrastructure efficiency, cloud architectures enable granular resource management impossible in traditional environments, with elastic scaling capabilities ensuring computational resources activate only when processing demands exist, eliminating the persistent energy consumption of idle infrastructure [1][2].

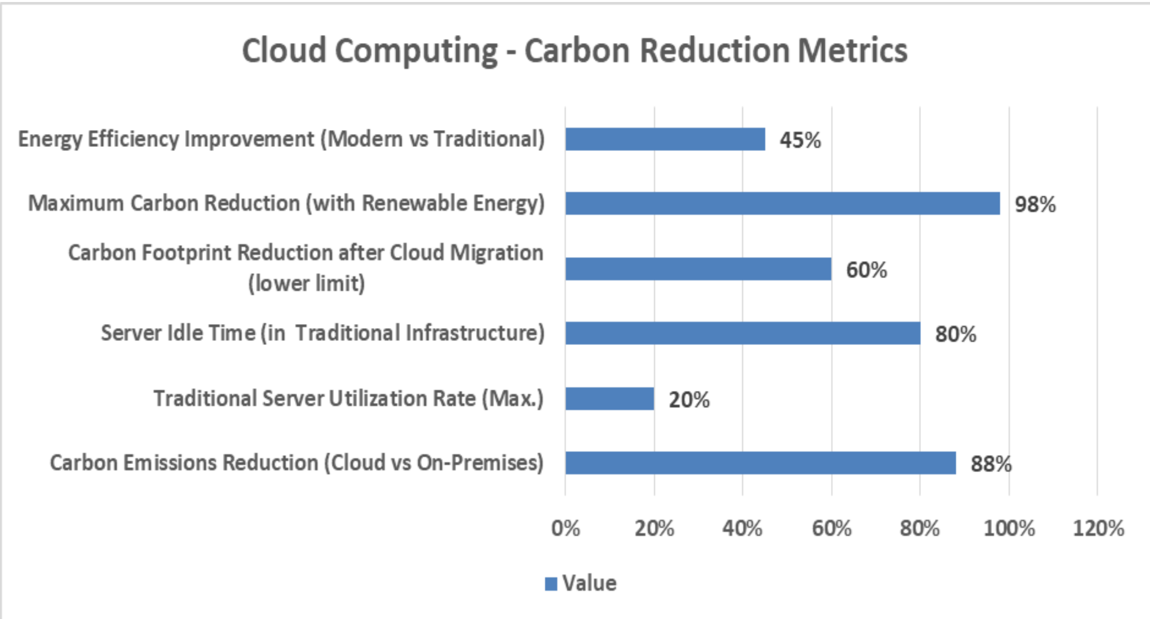

Figure 1: Cloud Computing - Carbon Reduction Metrics [1,2]

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

2. The Environmental Burden of Traditional Data Centers

Conventional in-premise data center infrastructure is a significant environmental problem in the insurance sector, exhibiting low resource utilisation, extremely high energy consumption patterns, and massive carbon emissions. For every unit of power that the computational processing uses, a typical data center allocates the rest of the 0.8 to 1.5 units to cooling, power distribution, and auxiliary systems. This wastage is explicable by outdated cooling technologies, improper handling of airflow, and electricity structures that had been built decades earlier, when the contemporary world's efficiency standards were invented. The insurance industry's reliance on legacy infrastructure exacerbates these problems; traditional facilities use around 40% more energy per computational unit compared to their modern hyperscale counterparts [3].

Server utilization in traditional enterprise data centers is often in the range of 10% to 20%, which implies that about 80-90% of the time, the computational resources are idle while consuming the baseline electrical power. This happens because the legacy infrastructure is not designed for dynamic provisioning and hence requires servers to be 'always-on' irrespective of the workload. Research indicates that traditional data centers waste between 60% and 70% of consumed energy through idle capacity and inefficient cooling systems [4]. Insurance carriers maintaining distributed data center networks face compounded inefficiencies, as smaller facilities cannot achieve the economies of scale available to hyperscale operations. Equipment refresh cycles generate substantial electronic waste streams, with decommissioned servers, storage arrays, and networking hardware containing hazardous materials requiring specialized disposal processes that contribute to environmental degradation [3].

Water consumption for cooling systems presents another critical environmental dimension. Traditional data centers employing evaporative cooling can consume between 3 and 5 million liters annually per megawatt of IT load, placing strain on local water resources, particularly in regions facing water scarcity. Cooling towers and computer room air conditioning units operate continuously, with water consumption scaling proportionally to computational load and ambient temperature conditions [4]. The carbon intensity of electricity sourcing further amplifies environmental impacts, as many legacy facilities draw power from regional grids heavily dependent on fossil fuel generation. Geographic constraints often positioned data centers in locations selected for business proximity rather than environmental considerations, resulting in operations in regions with carbon-intensive electrical grids or unfavorable climate conditions requiring intensive cooling [3].

Fixed capacity provisioning inherent to traditional infrastructure creates permanent environmental overhead. Insurance operations experience significant demand variability-peak periods during catastrophic events, renewal cycles, or regulatory reporting windows–yet infrastructure must remain sized for maximum anticipated load. Energy efficiency studies demonstrate that traditional data centers operate at optimal efficiency only 15-20% of the time, with the remaining operational periods characterized by resource underutilization and energy waste exceeding 45% of total consumption [4].

3. Strategic Approaches to Cloud Migration

Cloud migration strategies in insurance infrastructure go beyond mere rehosting exercises to basic architectural transformation that optimizes both operational performance and environmental impact. Traditional migrations based on "lift and shift" offer immediate benefits related to infrastructure consolidation, but cannot capture the full potential of sustainability that could be achieved by cloud-native architectures. It has been found from research that comprehensive architectural refactoring during cloud migration could reduce energy consumption by 40% to 60% compared to simple rehosting approaches, while carbon footprint reductions may reach up to 70% to 85% when combined with renewable energy sourcing and optimized workload management [5]. The transition of the

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/

Research Article

insurance sector requires careful strategic planning with a balanced approach to immediate operational needs and long-term objectives of sustainability [6].

Microservices architecture is basically a structural pattern of sustainable cloud operations. Breaking down these monolithic insurance programs, including the old policy administration software, claims processing systems, and underwriting engines, into highly discrete, independently deployable generalized services allows the allocation of resources more precisely in line with the realistic workload needs. The scaling of each microservice in a demand pattern is possible horizontally, meaning that only when the volumes of transactions warrant the provision of additional instances, they will be activated. This is in contrast to the monolithic architectures, where whole application stacks must remain up and running even when only portions of the provided functionality are in demand. The resource utilization rates of containerization technologies that support microservices deployments reach over 70%, compared to 10-20% typical for traditional virtual machine environments, which directly translates to reduced computational infrastructure and associated energy consumption. Insurance organizations adopting microservices architectures report reductions in operational costs ranging from 30% to 50%, while simultaneously improving energy efficiency between 45% and 60% [6].

The relationship between business functionality and environmental impact is further optimized with serverless computing paradigms. In Function-as-a-Service architectures, instantiation of execution environments occurs only for specific operations, totally eliminating persistent compute resources. Insurance workflows have events that drive processes: calculations of premiums triggered by policy changes, fraud detection algorithms triggered by claim submissions, or regulatory reporting generated on scheduled intervals. These kinds of insurance workflows naturally fit into serverless models. Studies have shown that serverless architecture can reduce energy consumption from 60% to 80% as compared to traditional always-on server infrastructure; some of the implementations achieve as high as 90% in low-demand periods [5]. Again, the billing is based on actual execution time measured on a millisecond scale, which means that organizations use resources only when they are processing active transactions [6].

The autoscaling feature inherent in cloud architectures enables dynamic management of resources that is not at all possible in fixed capacity environments. Horizontal pod autoscaling containers orchestrators keep track of application metrics like request rates, queue depths, or any other custom business indicator, and respond by automatically adjusting the number of replicas according to demand. Insurances where daily business cycles are predictable, where seasonal business cycles arise with periodic renewals, or where catastrophic events are unpredictable, change with infrastructures that enlarge and contract accordingly. A combination of horizontal and vertical scaling reduces average resource allocation by 50% to 65% relative to peak-provisioned static infrastructure with maintained performance service level agreements; this means energy savings ranging from 45% to 70% across operational cycles [5][6].

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

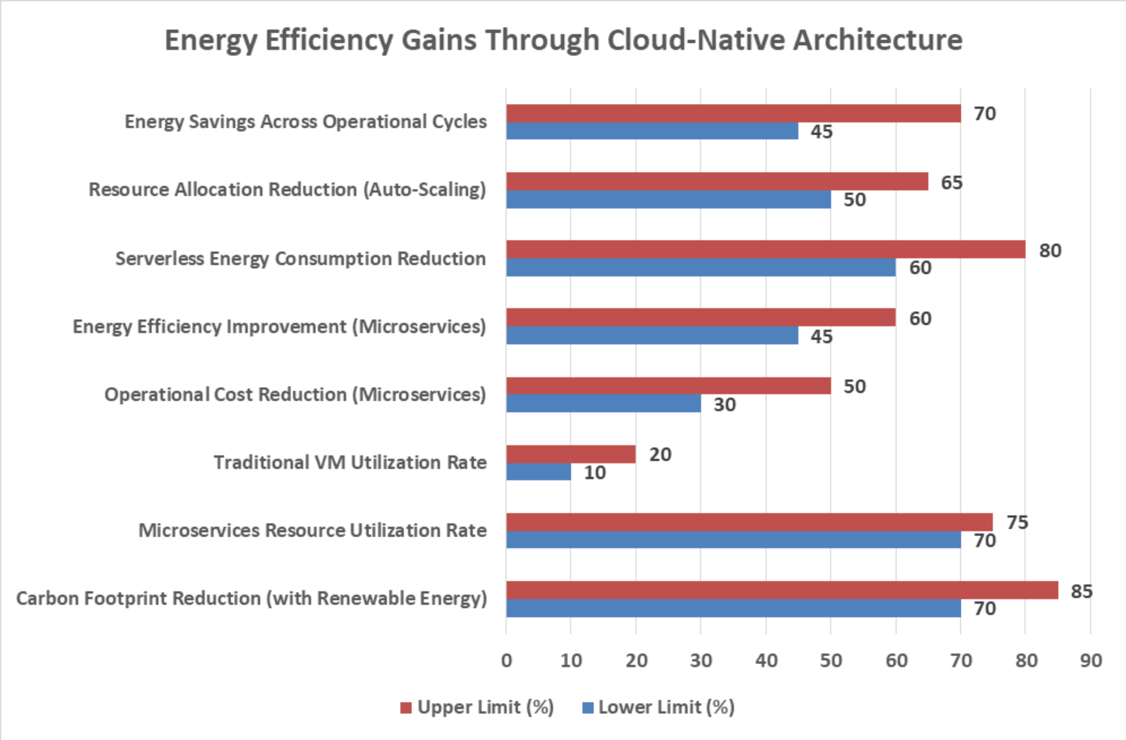

Figure 2: Energy Efficiency Gains Through Cloud-Native Architecture [5,6]

4. Technology-Enabled Environmental Management

Machine learning and artificial intelligence technologies have become key facilitators of environment optimization in cloud computing environments, making sustainability not periodic and manual but ongoing and automated. Anticipatory distribution of resources can be done by predictive software that examines past workload patterns, seasonal variations, and business cycle associations so that it can make proactive changes to distribution before demand tailgates appear. The accuracy of the state-of-the-art machine learning models that take into account the multi-dimensional data streams, such as transaction volumes, user behavior, and external triggers, is above 85 percent for demand prediction and enables infrastructure provisioning decision-making that reduces the risks of resource waste and performance degradation. However, the environmental cost of AI systems themselves requires careful consideration, with training large-scale models consuming between 502 and 1,287 kWh of energy, depending on model architecture and hardware configuration, equivalent to carbon emissions ranging from 200 to 500 kg CO2e per training cycle [7].

Another important area of intelligent environmental optimization is the automated management of the resource lifecycle. Cloud environments have accumulated orphaned resources - virtual machines that have been deployed to test but never decommissioned, storage volumes that have been detached due to instances that have been destroyed, yet remain costly in terms of resource usage and energy use, or load balancers that served an application that was destroyed. Machine learning classifiers that are trained to provide resource usage data, attachment status data, and access logs automatically recognize termination candidates with accuracy levels close to 92% allowing infrastructure hygiene to proceed automatically. Carbon footprint estimation models incorporating hardware specifications, power consumption profiles, and regional grid carbon intensity factors provide granular visibility into environmental impacts, with estimation accuracy achieving correlation coefficients above 0.95 when validated against measured power consumption data [8].

Intelligent job scheduling algorithms optimize batch processing workloads based on multiple environmental factors. Insurance operations generate substantial computational demands through actuarial modeling, regulatory reporting, financial close procedures, and data analytics pipelines—workloads exhibiting temporal flexibility regarding execution timing. Advanced scheduling systems consider grid carbon intensity forecasts, regional electricity pricing signals, and data center capacity utilization when determining optimal execution windows. Research demonstrates that carbon-aware scheduling can reduce workload-associated emissions by 30% to 40% compared to immediate execution strategies, with some implementations achieving 50% to 60% reductions by deferring processing to periods when renewable energy availability peaks or grid carbon intensity reaches daily minimums of 50-150 grams CO_2 per kWh compared to peak values of 400-800 grams CO_2 per kWh [8].

Configuration drift that reduces the efficiency of energy over time is detected through continuous monitoring and detection systems. ML classification as developing base performance and resource use patterns automatically alerts against deviations that can be exploited to optimize resource usage, identifying buggy behavior (memory leak that may cause the resources consumption to increase gradually) or resource wastefulness (code deployment that doubles or triples operations capacity to meet the same workload) or misconstrued auto-scaling rules (quantum underscaling that cannot scale-down when needed) by nature can trigger alarms with detailed metrics showing how much each metric has deviated and how that relates to past data. Detection sensitivity thresholds achieving precision rates of 88% to 92% enable automated alerting and self-healing remediation actions that restore optimal configurations, reducing operational carbon footprints by 15% to 25% through continuous optimization [7][8].

Optimization Capability	Performance Metric
AI Model Training Energy Consumption (kWh)	502-1,287
Carbon Emissions per Training Cycle (kg CO_2e)	200-500
Resource Termination Classification Accuracy	92%
Carbon Footprint Estimation Correlation	>0.95
Emission Reduction (Carbon-Aware Scheduling)	30-40%
Maximum Emission Reduction (Optimal Timing)	50-60%
Grid Carbon Intensity Minimum (grams CO_2/kWh)	50-150
Grid Carbon Intensity Peak (grams CO_2/kWh)	400-800

Table 1: AI-driven effectiveness and Carbon Footprint Reduction [7,8]

5. Measurement, Transparency, and Accountability

The best environmental management in cloud computing settings requires thorough measurement systems that go beyond IT's conventional measurement of performance to include both sustainability parameters at various organizational levels. The PUE ratio is a basic efficiency measure that describes the relationship between total facility energy consumption and energy delivered directly to the computing equipment. Hyperscale modern data centers can achieve PUE values ranging from 1.1 to 1.3, which means that for every unit of energy consumed for powering up servers, around 0.1 to 0.3 units are spent to cool down or support power distribution and auxiliary systems. On the other hand, through the integration of advanced security and monitoring mechanisms, sustainable cloud systems can report improvements in energy efficiency of 23% to 35% over traditional implementations while

2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/

preserving security performance metrics with intrusion detection accuracy rates of more than 94% and anomaly identification accuracy rates also above 94% [9].

The carbon accounting schemes classify emissions under three different scopes that are in accordance with the greenhouse gas protocols. Scope 1 emissions are those that occur at the sources of ownership or control, Scope 2 emission focuses on those that are a result of purchasing electricity, and Scope 3 emissions are those that are a result of purchasing hardware, manufacturing, and end-of-life disposal. Compression and pruning methods used to optimize machine learning models show potentially large reductions in various aspects of computational carbon footprints, with model compression being used to reduce model sizes by 60 to 80 percent without losing accuracy within 2 to 5 percent of initial performance. Pruning techniques removing redundant neural network parameters reduce computational requirements by 40% to 70%, directly translating to energy consumption reductions of 35% to 65% during inference operations [10].

Real-time monitoring dashboards provide granular visibility into environmental performance at workload, application, and business unit levels. These systems translate abstract carbon metrics into actionable operational intelligence, showing per-transaction carbon intensity and application-level energy consumption trends. Advanced platforms correlate computational workloads with carbon emissions, allowing organizations to attribute environmental impact to specific business activities. The location selection of data centers has a significant impact on carbon footprints, with regional variation in the carbon content in the grid varying between 50 grams CO_2 per kilowatt-hour in renewable-based grids and 800-900 grams CO_2 per kilowatt-hour in carbon-dependent grids. A quantized neural network implementation cuts memory bandwidth use by 75% and computational energy use by 50 to 60 percent compared to full-precision baseline implementations, which lowers an operational carbon footprint by 45 to 55 percent across the lifecycle of the model [10].

Integration of sustainability metrics into corporate governance structures ensures environmental considerations influence strategic decision-making. Sustainable cloud computing frameworks incorporating encryption, data mining, and continuous monitoring capabilities achieve system efficiency improvements of 28% to 42% while reducing overall energy consumption by 20% to 35% compared to conventional approaches [9]. IT scorecards incorporating carbon efficiency alongside availability, performance, and cost metrics embed environmental accountability into operational management, with quarterly board reporting routinely including environmental performance updates positioned alongside financial results and operational key performance indicators.

Measurement Category	Value/Range
Modern Hyperscale Data Center PUE	1.1-1.3
Energy Efficiency Improvement (Sustainable Systems)	23-35%
Machine Learning Model Compression Rate	60-80%
Accuracy Maintenance After Compression	2-5% loss
Computational Requirement Reduction (Pruning)	40-70%
Energy Consumption Reduction (Inference)	35-65%
Grid Carbon Intensity (Renewable Regions)	50 grams CO_2/kWh
Grid Carbon Intensity (Fossil Fuel Regions)	800-900 grams CO_2/kWh
System Efficiency Improvement (Integrated Framework)	28-42%

Table 2: Carbon accounting results [9, 10]

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

Conclusion

The combined essentialities of digital transformation and the need to become environmentally compliant have placed the insurance sector at a crossroads where there is a need to reconsider its perspective towards the technology infrastructure globally. The use of cloud computing goes beyond the historical questions of operational efficiency to have far-reaching environmental effects with energy efficiency through hyperscale, using renewable energies, and smarter resource management, showing significant carbon footprint benefits. The replacement of old fixed capacity infrastructure by a dynamically scalable cloud-native infrastructure not only removes steady state energy waste but also allows the computational needs to be calculated exactly to match the needs of the business. Machine learning and artificial intelligence capabilities are responsible for the fact that environmental management is no longer about intervention in environmental crises, but optimization of processes, constant detection of opportunities to increase efficiency, and automatic response measures. Holistic measurement systems reflecting a fine level of insight into environmental performance within organizational, application, and transaction boundaries empower the data-centered decision-making processes needed to realize ambitious decarbonization goals. Embedding sustainability metrics within the corporate governance frameworks will entrench accountability of the environment in the strategic technology decisions so that they are in resonance with the mounting regulatory demands and expectations of the stakeholders. The sustainability leadership of the insurance industry in using the cloud does not just limit itself to internal operational advantages but also extends into the sustainability of the wider technology ecosystem through its procurement standards, underwriting practices, and client risk consulting. Those organizations adopting holistic sustainable cloud plans show that environmental accountability and business agility, as well as resiliency and competitive advantage, do not conflict with one another, but lead towards successful emission cuts alongside vibrant biz agility, resiliency, and market positioning in a growing, progressively climate-conscious trade place.

References

[1] Meesam Raza et al., "Carbon footprint reduction in cloud computing: Best practices and emerging trends", International Journal of Cloud Computing and Database Management, 2024. [Online]. Available: https://www.computersciencejournals.com/ijccdm/article/58/5-1-7-236.pdf

[2] Chandra Prakash Kathroju, "The Evolution of Cloud Adoption in Insurance: Guidewire's Strategic Vision", European Journal of Computer Science and Information Technology, Jun. 2025. [Online]. Available: https://eajournals.org/wp-content/uploads/sites/21/2025/06/The-Evolution-of-Cloud-Adoption.pdf

[3] European Insurance and Occupational Pensions Authority, "The Future of Cloud Computing and AI in the EU Insurance Sector", 2024. [Online]. Available: https://www.eiopa.europa.eu/publications/future-cloud-computing-and-ai-eu-insurance-sector_en

[4] Salil Bharany et al., "A Systematic Survey on Energy-Efficient Techniques in Sustainable Cloud Computing", MDPI, 2022. [Online]. Available: https://www.mdpi.com/2071-1050/14/10/6256

[5] C.A. Silva et al., "A review on the decarbonization of high-performance computing centers", ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032123008778

[6] Farheem Shakeel, "Cloud transformation in insurance: Modernizing processes for the digital age", Damco Group, Jul. 2025. [Online]. Available: https://www.damcogroup.com/blogs/cloud-transformation-in-insurance-for-modernizing-processes

Journal of Information Systems Engineering and Management
2025, 10(62s)
e-ISSN: 2468-4376
https://jisem-journal.com/
Research Article

[7] Sophia Falk et al., "More than Carbon: Cradle-to-Grave environmental impacts of GenAI training on the Nvidia A100 GPU", arXiv, Aug. 2025. [Online]. Available: https://arxiv.org/html/2509.00093v1

[8] Baolin Li et al., "Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems", arXiv, 2023. [Online]. Available: https://arxiv.org/pdf/2306.13177

[9] Qian He and Hong He, "A Novel Method to Enhance Sustainable Systems Security in Cloud Computing Based on the Combination of Encryption and Data Mining", MDPI, 2020. [Online]. Available: https://www.mdpi.com/2071-1050/13/1/101

[10] Stow, May et al., "Reducing Carbon Footprint of Machine Learning Through Model Compression and Pruning", IJISRT, Aug. 2025. [Online]. Available: https://www.ijisrt.com/assets/upload/files/IJISRT25AUG970.pdf