Jurnal Komputer

Vol. 1, No. 1, October 2026

SAH PUBLISHER

Data Center Digital Twin: Real-Time 3D Visualization and AI-Based Expansion Recommendation

Muhammad Luqman Al-Hakim[1*], Mirko Prasetyo[1], Amri Khoirudin[1]

[1]Universitas Sugeng Hartono, Sukoharjo, Indonesia

Email: luqman.alhakim003@gmail.com

Abstract. The rapid growth of global data continuously pushes data center infrastructure toward its physical limits, yet deciding when and how to expand capacity remains a critical challenge for operators. Sudden traffic surges can exceed the throughput of physical infrastructure, while uncontrolled power consumption and overheating threaten densely packed server racks. This paper presents a data center digital twin that unifies real-time telemetry, interactive 3D visualization, and an AI-based expansion recommendation engine in a single web-based system. The architecture follows an eight-block digital twin design: physical facilities, IoT data acquisition, data integration, the virtual twin state, prediction and explainability, optimization and feedback, the user dashboard, and a closed feedback loop. The twin state is maintained as a structured vector covering power infrastructure, cooling, environment, workload distribution, server utilization, network health, and equipment health. An expansion scoring engine computes a vertical score for specification upgrades and a horizontal score for adding nodes or racks from utilization, traffic, power, and cost inputs, and derives a confidence score from the gap between the two. The 3D scene, rendered with WebGL, animates data flow as dynamic particles and maps temperature and load telemetry onto rack heat colors, while a control panel with sliders and quick scenario presets enables zero-risk what-if simulation. A functional prototype with sixteen racks and one hundred twenty-eight servers demonstrates centralized visual monitoring, proactive risk testing through quick scenario presets, and data-driven expansion recommendations with confidence scores. The main contribution is an integrated digital twin architecture that couples real-time 3D visualization with a transparent scoring-based expansion decision engine for data center operators.

Received August 2026 / Revised September 2026 / Accepted October 2026

This work is licensed under a Creative Commons Attribution 4.0 International License.



Keywords: digital twin; data center; capacity planning; 3D visualization; what-if simulation; energy efficiency

INTRODUCTION

Global data volume continues to grow at an unprecedented rate, and data centers form the physical backbone of that growth. Modern facilities concentrate thousands of servers, storage systems, and network devices in dense racks, consuming large amounts of electricity and generating substantial heat [1]. Operating such infrastructure efficiently requires continuous awareness of power, thermal, network, and utilization conditions across thousands of sensor points. A digital twin, a virtual representation of a physical system synchronized with real-time data, provides a natural framework for this task [2][3]. Recent work has demonstrated digital twins for data centers covering energy management, thermal behavior, and network operations [4][5][6][7].

Three persistent challenges motivate this work. First, sudden traffic surges can exceed the throughput capacity of physical infrastructure, creating network bottlenecks that are difficult to anticipate [8]. Second, uncontrolled power consumption and overheating in dense racks risk hardware damage and service degradation, making thermal management a first-order concern [9][10][11][12]. Third, capacity expansion decisions are typically based on static reports and operator intuition, although both vertical scaling (upgrading server specifications) and horizontal scaling (adding nodes or racks) carry significant cost and risk [13].

Existing approaches address parts of this problem. Data center infrastructure management (DCIM) platforms provide monitoring dashboards but rarely couple them with predictive expansion logic. Simulation toolkits such as CloudSim and GreenCloud support what-if analysis of resource provisioning and energy-aware scheduling [14][15], but they operate offline rather than as live synchronized twins. Thermal studies rely on computational fluid dynamics and sensor placement optimization [16][17], and

energy research targets power usage effectiveness (PUE) through workload scheduling and cooling optimization [18][19][20], while renewable integration and carbon-aware operation add further dimensions to the efficiency problem [21][22][23]. What is missing is a single interactive environment where operators can observe the live state of the facility in 3D, test expansion scenarios without risk, and receive transparent expansion recommendations. This paper presents such a system: a web-based data center digital twin with real-time 3D visualization, a simulation control panel, and a scoring-based expansion recommendation engine. The main contributions are (1) an eight-block digital twin architecture for data centers with an explicit twin state vector, (2) a transparent vertical-versus-horizontal expansion scoring engine with a confidence measure, and (3) a functional prototype demonstrating centralized 3D monitoring and zero-risk what-if simulation.

METHODS

System Architecture

The proposed system follows the architecture illustrated in Fig. 1, organized into eight functional blocks. The physical facilities block covers the data center floor, power distribution units, and uninterruptible power supplies, producing facility resources such as power, cooling capacity, and space. The IoT and data acquisition block collects measurements from four sensor groups: environmental sensors (temperature, humidity, airflow), power monitoring (rack-level power metering, PDU load), server monitoring (CPU, RAM, load), and cooling monitoring (CRAC status, setpoints), with liquid cooling increasingly adopted for high-density racks [24]. The data integration layer performs data ingestion from IoT devices, server logs, and building management records, followed by time synchronization, data cleaning, and feature engineering into a unified data repository. The digital twin core maintains the virtual twin state S(t) = {F(t), C(t), E(t), W(t), P(t), N(t), M(t)}, representing infrastructure and power facilities, the cooling system, environmental conditions, workload distribution, server utilization, network and bandwidth health, and equipment health and maintenance; the state is updated continuously from new data. The prediction and explainability block applies predictive models to metrics such as failure and PUE, estimates performance, and uses explainable AI to identify drivers of hotspot formation or cooling inefficiency. The optimization and feedback block searches for optimal setpoints and workload placements and issues decisions and recommendations. The user interface and dashboard block presents real-time data center status, the thermal floor view, twin status, PUE trends, hotspot alerts, and a recommendation viewer. Finally, a closed feedback loop implements recommended setpoints, observes the real facility response, updates the twin state, and improves subsequent recommendations [25][26].

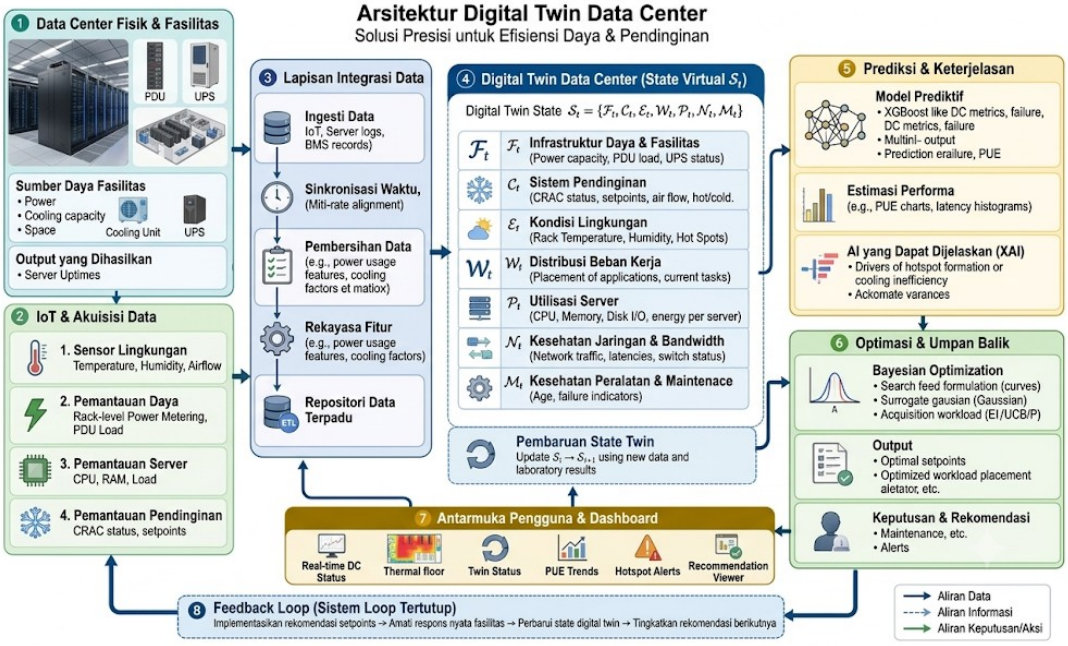


Figure 1. Data Center Digital Twin Architecture

Data Acquisition and Integration

Telemetry acquisition targets the four physical quantities that dominate data center risk: power, heat, network load, and compute utilization. Rack-level power metering captures per-rack consumption, complementing PDU load readings at the distribution level. Environmental probes measure inlet temperature, humidity, and airflow, quantities governed by the thermal guidelines for data processing environments [27]. Server agents report CPU, memory, disk I/O, and per-server energy, while CRAC units report status and setpoints. Effective use of this telemetry depends on sensor placement and calibration quality; prior work has shown that intelligent sensor placement improves hot server detection [16], and knowledge-based calibration keeps surrogate models aligned with physical behavior [25]. The integration layer aligns multi-rate streams in time, removes missing values and outliers, and engineers features such as power usage factors and cooling factors before storing them in the unified repository [6].

Expansion Recommendation Engine

The decision intelligence of the system is a scoring engine that evaluates two expansion strategies, illustrated in Fig. 2. Vertical expansion increases the specifications of existing servers (CPU and RAM) and is indicated when compute utilization is persistently high. Horizontal expansion adds new nodes or racks and is indicated when data and network load, power draw, or cost pressure dominate. Each strategy receives a cumulative score from conditional point additions listed in Table 1. The vertical score V accumulates 30 points when utilization exceeds 85% and 20 additional points when it exceeds 95%. The horizontal score H accumulates 25 points when network traffic exceeds 12 Gbps, 15 points when power draw exceeds 1000 kW, and 18 points when monthly server cost exceeds $100,000. The strategy with the higher score becomes the recommendation; when both scores are zero, the system reports that capacity is sufficient. A confidence score, capped at 96%, is derived from the gap between the two scores: the larger the difference, the higher the confidence in the recommendation. Threshold rules complement the scoring by triggering warning states automatically when compute utilization crosses 65% or network capacity approaches its safe limit. Unlike black-box predictors, this engine is fully transparent: every point contribution can be traced to a measured quantity and a stated rule, which supports explainable operational decisions [28].

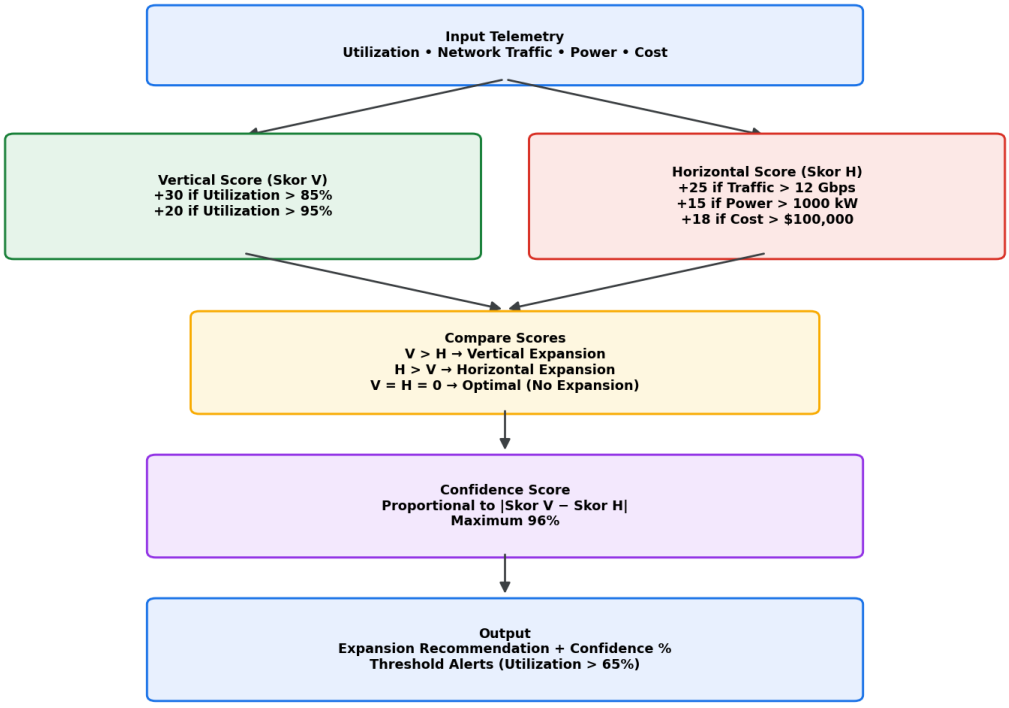


Figure 2. Expansion Scoring Algorithm of the Recommendation Engine

Table 1. Expansion Scoring Rules

Strategy	Condition	Points
Vertical (Skor V)	Utilization > 85%	+30
Vertical (Skor V)	Utilization > 95%	+20
Horizontal (Skor H)	Network traffic > 12 Gbps	+25
Horizontal (Skor H)	Power draw > 1000 kW	+15
Horizontal (Skor H)	Monthly server cost > $100,000	+18

Interactive 3D Visualization

The spatiotemporal visual layer renders the data center floor in 3D using WebGL through the Three.js library, running directly in the browser with orbit, zoom, and pan camera navigation. Data flow is visualized as dynamic pulse particles whose intensity reflects the volume of active traffic entering or leaving the facility. Each server rack changes color dynamically based on temperature and performance load telemetry, producing a live heat map of the floor that makes hotspots immediately visible [29]. Below the 3D scene, metric cards present key indicators with status color codes and mini-charts showing recent history. The interactive control center on the right provides dynamic input sliders for operational variables (power consumption, network traffic, inbound and outbound data, performance utilization, and monthly server cost), quick scenario preset buttons ranging from safe conditions (Startup) to emergency conditions (Overload), and the expansion recommendation panel. Because all scenario changes apply to the simulated twin rather than the physical facility, operators can test the infrastructure response to load surges before applying real changes, realizing zero-risk what-if analysis [14].

RESULT AND DISCUSSION

Prototype Implementation

The functional prototype implements the full interaction concept, as shown in Fig. 3. The 3D scene renders the server floor with 16 racks and 128 servers online, with data flow particles animating across the floor. The control panel sliders adjust power consumption (500 kW), network traffic (3.0 Gbps), inbound data (20 TB/day), outbound data (15 TB/day), performance utilization (55%), and monthly server cost ($35,000) in real time, and the quick scenario buttons (Startup, Menengah, Enterprise, Kritis, Overload) apply complete load profiles instantly. The metric cards at the bottom display each indicator with its current value and mini-chart history, color-coded by status. The expansion recommendation panel shows the verdict "OPTIMAL — Kapasitas Mencukupi" with a 95% confidence score for the demonstrated configuration, alongside the vertical and horizontal score bars. The prototype confirms that real-time telemetry, 3D visualization, and decision support can coexist in one operator view, and that scenario testing produces immediate visual and numerical feedback without touching physical hardware.

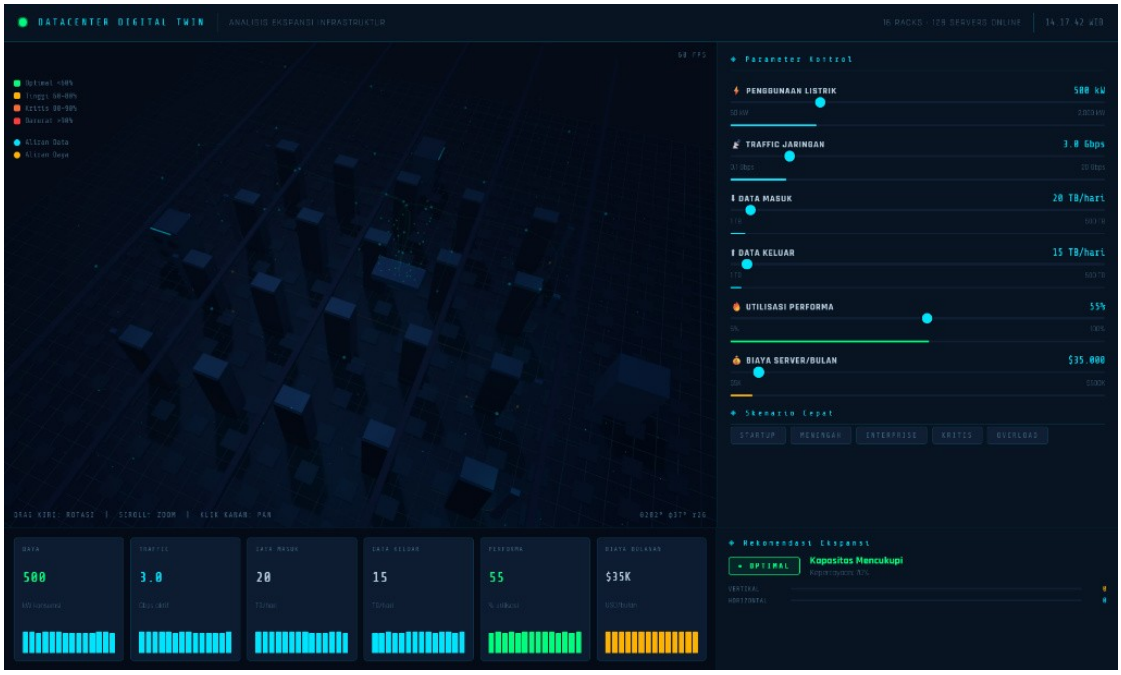


Figure 3. Implemented Dashboard of the Data Center Digital Twin Prototype

Expansion Scoring and Decision Logic

The scoring engine was exercised across the scenario presets to verify its decision behavior. Under low-load scenarios, both scores remain at zero and the system correctly reports sufficient capacity. As utilization is raised past 85%, the vertical score accumulates and the recommendation switches to vertical expansion with rising confidence; conversely, raising network traffic past 12 Gbps or power past 1000 kW drives the horizontal score and flips the recommendation to horizontal expansion. The confidence score behaves as designed, approaching its 96% cap when one strategy dominates and dropping toward zero when the scores are close, which honestly signals ambiguous cases to the operator rather than forcing a recommendation. Compared with simulation-only approaches such as CloudSim and GreenCloud [14] [15], the engine operates on live twin state and presents its reasoning transparently. Compared with predictive models for workload and energy [30][31][18], the rule-based scoring trades predictive depth for full explainability, which is appropriate for an operational dashboard where every recommendation must be auditable. The main limitation is that the prototype evaluates scoring logic on simulated telemetry rather than on measurements from a production facility; field validation with real sensor streams and historical expansion decisions remains future work.

CONCLUSION

This paper presented a data center digital twin that integrates real-time telemetry, interactive 3D visualization, and an AI-based expansion recommendation engine. The eight-block architecture maintains an explicit twin state vector spanning power, cooling, environment, workload, utilization, network, and equipment health. The scoring engine transparently compares vertical and horizontal expansion strategies and reports a confidence score for each recommendation. The prototype demonstrates centralized visual monitoring through a 3D floor view with live heat mapping, zero-risk what-if simulation through a control panel with sliders and scenario presets, and data-driven expansion decisions through the recommendation panel. Future work includes connecting the twin to live sensor streams, incorporating predictive models for failure and PUE alongside the rule-based scoring, and validating recommendations against historical capacity planning decisions. Digital-twin-driven real-time optimization has likewise been demonstrated for smart factory production systems using edge computing and Industrial IoT [32], while explainable AI frameworks for fault diagnosis and remaining-useful-life prediction offer a complementary direction for the predictive models envisioned as future work [33].

REFERENCES

[1] M. Dayarathna, Y. Wen, and R. Fan, "Data Center Energy Consumption Modeling: A Survey," IEEE Commun. Surveys Tuts., vol. 18, no. 1, pp. 732-794, 2016. DOI: 10.1109/COMST.2015.2481183

[2] J. Athavale, C. Bash, W. Brewer, M. Maiterth, D. Milojicic, H. Petty, and S. Sarkar, "Digital Twins for Data Centers," Computer, vol. 57, no. 10, pp. 151-158, 2024. DOI: 10.1109/MC.2024.3436945

[3] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital Twin: Enabling Technologies, Challenges and Open Research," IEEE Access, vol. 8, pp. 108952-108971, 2020. DOI: 10.1109/ACCESS.2020.2998358

[4] W. Brewer, M. Maiterth, V. Kumar, R. Wojda, S. Bouknight, J. Hines, W. Shin, S. Greenwood, D. Grant, W. Williams, and F. Wang, "A Digital Twin Framework for Liquid-Cooled Supercomputers as Demonstrated at Exascale," in Proc. Int. Conf. for High Performance Computing, Networking, Storage and Analysis (SC '24), Atlanta, GA, USA, 2024, pp. 1-18. DOI: 10.1109/SC41406.2024.00029

[5] T. I. Zohdi, "A Digital-Twin and Machine-Learning Framework for Precise Heat and Energy Management of Data-Centers," Computational Mechanics, vol. 69, no. 6, pp. 1501-1516, 2022. DOI: 10.1007/s00466-022-02152-3

[6] Z. Zhang, Y. Zeng, H. Liu, C. Zhao, F. Wang, and Y. Chen, "Smart DC: An AI and Digital Twin-Based Energy-Saving Solution for Data Centers," in Proc. NOMS 2022 - IEEE/IFIP Network Operations and Management Symposium, Budapest, Hungary, 2022, pp. 1-6. DOI: 10.1109/NOMS54207.2022.9789853

[7] H. Hong, Q. Wu, F. Dong, W. Song, R. Sun, T. Han, C. Zhou, and H. Yang, "NetGraph: An Intelligent Operated Digital Twin Platform for Data Center Networks," in Proc. ACM SIGCOMM 2021 Workshop on Network-Application Integration (NAI '21), Virtual Event, 2021, pp. 26-32. DOI: 10.1145/3472727.3472802

[8] W. Zhang, Y. Wen, Y. W. Wong, K. C. Toh, and C.-H. Chen, "Towards Joint Optimization over ICT and Cooling Systems in Data Centre: A Survey," IEEE Commun. Surv. Tutorials, vol. 18, no. 3, pp. 1596-1616, 2016. DOI: 10.1109/COMST.2016.2545109

[9] Y. Li, C. Zhu, X. Li, and B. Yang, "A Review of Non-Uniform Load Distribution and Solutions in Data Centers: Micro-Scale Liquid Cooling and Large-Scale Air Cooling," Energies, vol. 18, no. 1, 149, 2025. DOI: 10.3390/en18010149

[10] R. Zhao, Y. Du, X. Yang, Z. Zhou, W. Wang, and X. Yang, "A Critical Review on the Thermal Management of Data Center for Local Hotspot Elimination," Energy and Buildings, vol. 297, 113486, 2023. DOI: 10.1016/j.enbuild.2023.113486

[11] Y. Du, Z. Zhou, X. Yang, X. Yang, C. Wang, J. Liu, and J. Yuan, "Dynamic Thermal Environment Management Technologies for Data Center: A Review," Renewable and Sustainable Energy Reviews, vol. 187, 113761, 2023.

[12] S. Xu, H. Zhang, and Z. Wang, "Thermal Management and Energy Consumption in Air, Liquid, and Free Cooling Systems for Data Centers: A Review," Energies, vol. 16, 1279, 2023. DOI: 10.3390/en16031279

[13] E. Cortez, A. Bonde, A. Muzio, M. Russinovich, M. Fontoura, and R. Bianchini, "Resource Central: Understanding and Predicting Workloads for Improved Resource Management in Large Cloud Platforms," in Proc. 26th ACM Symp. Operating Systems Principles (SOSP), Shanghai, China, 2017, 15 pp. DOI: 10.1145/3132747.3132772

[14] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, "CloudSim: A Toolkit for Modeling and Simulation of Cloud Computing Environments and Evaluation of Resource Provisioning Algorithms," Software: Practice and Experience, vol. 41, no. 1, pp. 23-50, 2011. DOI: 10.1002/spe.995

[15] D. Kliazovich, P. Bouvry, and S. U. Khan, "GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers," J. Supercomput., vol. 62, no. 3, pp. 1263-1283, 2012. DOI: 10.1007/s11227-010-0504-1

[16] X. Wang, X. Wang, G. Xing, J. Chen, C.-X. Lin, and Y. Chen, "Intelligent Sensor Placement for Hot Server Detection in Data Centers," IEEE Trans. Parallel Distrib. Syst., vol. 24, no. 8, pp. 1577-1588, 2013. DOI: 10.1109/TPDS.2012.254

[17] W.-X. Chu and C.-C. Wang, "CFD Investigation of Airflow Management in a Small Container Data Center," IEEE Trans. Compon. Packag. Manuf. Technol., vol. 9, no. 11, pp. 2177-2188, 2019. DOI: 10.1109/TCPMT.2019.2916017

[18] Y. Ran, H. Hu, Y. Wen, and X. Zhou, "Optimizing Energy Efficiency for Data Center via Parameterized Deep Reinforcement Learning," IEEE Trans. Serv. Comput., vol. 16, no. 2, pp. 1310-1323, 2023. DOI: 10.1109/TSC.2022.3184835

[19] A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-Aware Resource Allocation Heuristics for Efficient Management of Data Centers for Cloud Computing," Future Gener. Comput. Syst., vol. 28, no. 5, pp. 755-768, 2012.

[20] Y. Li, Y. Wen, D. Tao, and K. Guan, "Transforming Cooling Optimization for Green Data Center via Deep Reinforcement Learning," IEEE Trans. Cybern., vol. 50, no. 5, pp. 2002-2013, 2020. DOI: 10.1109/TCYB.2019.2927410

[21] I. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, "Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy," in Proc. ACM Int. Conf. Architectural Support for Programming Languages and Operating Systems (ASPLOS), Houston, TX, USA, 2013, pp. 51-64.

[22] M. Xu and R. Buyya, "Managing Renewable Energy and Carbon Footprint in Multi-Cloud Computing Environments," J. Parallel Distrib. Comput., vol. 135, pp. 191-202, 2020. DOI: 10.1016/j.jpdc.2019.09.015

[23] M. Avgerinou, P. Bertoldi, and L. Castellazzi, "Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency," Energies, vol. 10, no. 10, 1470, 2017. DOI: 10.3390/en10101470

[24] M. Azarifar, M. Arik, and J.-Y. Chang, "Liquid Cooling of Data Centers: A Necessity Facing Challenges," Applied Thermal Engineering, vol. 247, 123112, 2024. DOI: 10.1016/j.applthermaleng.2024.123112

[25] R. Wang, X. Zhou, L. Dong, Y. Wen, R. Tan, L. Chen, G. Wang, and F. Zeng, "Kalibre: Knowledge-based Neural Surrogate Model Calibration for Data Center Digital Twins," in Proc. 7th ACM Int. Conf. on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '20), Virtual Event, Japan, 2020, pp. 200-209. DOI: 10.1145/3408308.3427982

[26] Z. Cao, R. Wang, X. Zhou, and Y. Wen, "Reducio: Model Reduction for Data Center Predictive Digital Twins via Physics-Guided Machine Learning," in Proc. 9th ACM Int. Conf. on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '22), Boston, MA, USA, 2022, pp. 1-10. DOI: 10.1145/3563357.3564050

[27] ASHRAE Technical Committee 9.9, Thermal Guidelines for Data Processing Environments, 5th ed. Atlanta, GA, USA: ASHRAE, 2021.

[28] H. Yuan, Z. Zhang, D. Yang, T. Xue, D. Wen, and G. Yao, "An Intelligent Thermal Management Strategy for a Data Center Prototype Based on Digital Twin Technology," Applied Sciences, vol. 15, art. 7675, 2025. DOI: 10.3390/app15147675

[29] B. Bergeron, M. Hubbell, D. Mojica, A. Reuther, W. Arcand, D. Bestor, D. Burrill, C. Byun, V. Gadepally, M. Houle, H. Jananthan, M. Jones, P. Luszczek, P. Michaleas, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and J. Kepner, "Supercomputer 3D Digital Twin for User Focused Real-Time Monitoring," in Proc. IEEE High Performance Extreme Computing Conf. (HPEC), 2024, pp. 1-8.

[30] D. Saxena, J. Kumar, A. K. Singh, and S. Schmid, "Performance Analysis of Machine Learning Centered Workload Prediction Models for Cloud," IEEE Trans. Parallel Distrib. Syst., 2023. DOI: 10.1109/TPDS.2023.3240567

[31] D. Saxena and A. K. Singh, "Workload Pattern Learning-Based Cloud Resource Management Models: Concepts and Meta-Analysis," IEEE Trans. Sustainable Comput., vol. 10, no. 3, pp. 418-438, 2025. DOI: 10.1109/TSUSC.2024.3456429

[32] S. Suyahman, D. U. Iswavigra, H. Wibowo, A. B. Trisnawan, A. Wicaksono, et al., "Digital Twin Driven Real Time Performance Optimization of Smart Factory Production Systems Using Edge Computing and Industrial Internet of Things Architecture," Int. J. Industrial Innovation and Mechanical Engineering, vol. 1, no. 2, 2024.

[33] S. Suyahman, D. Prasetyo, A. B. Trisnawan, A. Wicaksono, and M. Furqon, "Explainable Artificial Intelligence Framework for Interpretable Fault Diagnosis and Remaining Useful Life Prediction in Smart Industrial Rotating Machinery," IJMICSE: Int. J. Mechanical, Industrial and Control Systems Engineering, 2024.