Received: 7 October 2022 | Revised: 5 January 2023 | Accepted: 27 January 2023
DOI: 10.1002/cpe.7658

RESEARCH ARTICLE

WILEY

A cyber-physical management system for medium-scale solar-powered data centers

Marios D. Dikaiakos[1] | Nikolas G. Chatzigeorgiou[2] | Athanasios Tryfonos[1] | Andreas Andreou[1] | Nicholas Loulloudes[1] | George Pallis[1] | George E. Georgiou[2]

[1]Department of Computer Science, University of Cyprus, Nicosia, Cyprus

[2]PV Technology Lab, Department of Electrical & Computer Engineering, University of Cyprus, Nicosia, Cyprus

Correspondence
Marios D. Dikaiakos, Department of Computer Science, University of Cyprus, Nicosia, Cyprus.
Email: mdd@ucy.ac.cy

Funding information
This work was supported by the ENEDI Project, which was co-financed by the European Union (ERDF) and national funds of Greece and Cyprus (Cross-Border Cooperation Programme Greece Cyprus 2014 - 2020)

Summary

The effort to reduce the environmental impact and carbon footprint of data-center operations has led to the emergence of "green" data centers, which are designed to reduce energy consumption and to increase their use of Renewable Energy Sources (RES). Despite the advances demonstrated by hyper-scale facilities in energy efficiency and the use of green energy, small and medium-scale data centers, which contribute to over 50% of the total electricity consumption and carbon emissions of the sector, face significant challenges in the adoption and exploitation of RES. In this article, we present the steps taken to transform a medium-scale, academic data center into a "green" one that uses solar power. In particular, we describe the design and implementation of: (i) a collocated photovoltaic facility and (ii) a cyber-physical system comprising IoT sensor devices, a microservices platform, and a visualization and analytics dashboard that supports the configuration and monitoring of the infrastructure. Using data collected from the platform and dashboard, we show the environmental and financial advantages derived from this transformation, and the potential that arises from the availability of integrated operational data.

KEYWORDS

containers, "green" data centers, IoT sensors, microservices, photovoltaics, self-consumption

1 | INTRODUCTION

Cloud computing has become a general-purpose technology that provides the underlying IT infrastructure across numerous sectors of human activity. The rapid growth of cloud services was made possible thanks to a concomitant expansion in data center deployments worldwide.[1] Industry analysts estimate that, between 2015 and 2020, the expansion of data centers was able to sustain a 3-fold increase in computing workloads and IP traffic, and a 5-fold increase in storage capacity.[2] Public cloud services are typically provided through the utility computing model,[3] wherein access to cloud resources is charged proportionally to their actual use.[4,5] This model offers financial savings to cloud users who can thereby reduce their overall capital and operational IT expenditures, budgeting them according to evolving computing needs. The utility computing model enables also cloud operators to capitalize on the multiplexing of own resources across many customers, to amortize maintenance and IT management costs across thousands of servers, and to adopt automated techniques for data center management.[1]

Energy is one of the main factors that drive data center operational expenses,[1] total cost of ownership (TCO),[6,7] and, consequently, the pricing of cloud services and profit margins of cloud operators. The expansion of data center deployments around the world has resulted in data center energy consumption rising from an estimated 153 terawatt-hours (TWh) in 2005, to 194 TWh in 2010, and 205 TWh in 2018.[8,9] Earlier estimates predicted that data center energy consumption would reach up to 13% of total electricity use worldwide by 2030.[10-12] These predictions did not materialize thanks to rapid improvements in the energy efficiency of computer hardware[8,13] and data center facilities (power distribution, cooling, etc.),[14-16] particularly in the hyper-scale data centers of large operators like Amazon, Google and Microsoft,[1,17] which have achieved near-optimal Power Usage Effectiveness (PUE) values below 1.1–1.2.[1,18-21] Also, due to the financial crisis of 2008 and the ensuing worldwide economic slowdown, which, by 2010, led to a significant reduction in actual servers installed, compared to forecasts.[9]

According to more recent estimates (2019, 2020), the electricity consumption of data centers has reached 1.1%–1.5% of total electricity use worldwide,[8] exceeding 2% in the USA[22] and growing at a 4.3% rate per annum globally.[23] This rate is expected to increase due to the accelerated pace of digitalization triggered by the Covid-19 pandemic[22,24] and the rapidly growing popularity of energy-demanding applications in machine learning[25] and crypto-currency mining.[26,27]

Nevertheless, according to estimates published by the International Energy Agency in 2019, the worldwide energy demand of traditional and cloud data centers (nonhyper-scale) dropped from 60.9 and 76.3 TWh in 2018 to 32.6 and 71.6 TWh in 2021, respectively, whereas the energy demand of hyper-scale data centers increased during the same period from 60.9 to 86.6 TWh.[28] These figures suggest that despite the expansion of data centers worldwide, the shift of investments and workloads from small and medium facilities to hyper-scale data centers with better PUE has led to a small reduction of their overall energy demand: from 197.8 TWh in 2018 to 190.8 TWh in 2021.

This trend does not necessarily hold true across all regions and countries, where the rate and impact of data-centers' expansion are determined by local characteristics in market size, needs, cloud service penetration, and regulation. For instance, in the EU, the energy consumption of data centers in 2018 reached 76.8 TWh (2.7% of total EU electricity demand) and was expected to rise to 98,52 TWh (3.21% of total electricity) by 2030.[29] In smaller countries with a dynamic data-center sector, such as Denmark and Ireland, data-center electricity is expected to grow rapidly and occupy a significant share of overall electricity consumption, reaching by 2028 a 15% and 30% of total electricity demand, respectively.[29] Finally, the development of edge computing services is expected to push energy consumption of data centers further, with edge data-center energy consumption in the EU expected to rise from 2% to 12% between 2018 and 2025.[29]

However, electricity production relies predominately on carbon-intensive fossil fuels and brings a substantial toll on our environment. To address this problem, governments around the world adopt renewable portfolio standards, set goals that aim at reducing electric power carbon emissions by providing incentives for expanding the production and use of energy from renewable energy sources (RES).[30,31] As data-center deployment accelerates and expands geographically,[32] the contribution of data centers to the worldwide carbon footprint increases and brings a substantial toll on our environment.[16,22,33,34] Therefore, leading operators of hyper-scale data centers like Google,[35] Amazon[36] and Microsoft[21,37] have undertaken initiatives to establish "green" data centers, which are designed to be powered up to 100% with clean and renewable energy, in a quest to mitigate carbon emissions and reduce their energy costs.

Nonetheless, more than 50% of overall data center energy consumption still comes from small and medium data centers,[17,23] whose energy-efficiency is typically inferior to that of large and hyper-scale facilities. With the advent of 5G[38] and multi-cloud edge computing,[39] and the proliferation of edge and fog services,[40] the deployment of small, medium, micro, and nano data centers is expected to expand rapidly inside urban areas,[38,41] in order to cope with the growing data processing requirements that arise from smart city, connected vehicles, smart manufacturing and other cyber-physical applications. More often than not, small and medium data centers are developed and upgraded incrementally, and their operators follow procurement and installation practices different than those applied in hyper-scale facilities. Typically, small and medium data centers use diverse equipment for computing, cooling, power supply, and energy management. Consequently, as indicated by surveys and studies, small and medium data centers waste more energy than hyper-scale facilities, exhibiting average PUE values between 1.6–2.[1,34,42] Also, the management of their operations cannot be easily streamlined, automated, or federated as it relies upon different subsystems, which are typically proprietary, use silo-ed data and incompatible analytics and interfaces.

Therefore, when operators of small or medium-size data centers wish to improve the energy efficiency of their facilities, to reduce their carbon footprint, to adapt their cloud operations and pricing policies to a dynamically changing mixture of energy sources, and to engage in resource sharing with other data centers, they are faced with numerous challenges including:

- Technical, regulatory, and financial hurdles in choosing, procuring, installing, and operating renewable energy production facilities.
- The integration of diverse subsystems, hardware sensors and software agents that collect data about the status of data-center resources, in order to make operational, energy, environmental and financial-related data available through a common platform that would facilitate processing, actionable analytics, and visualization.
- The configuration, maintenance and continuous upgrade of various software components of the operational platform.
- The capability to intelligently adjust data center operations and service pricing to variations of renewable energy production, changing environmental conditions, and fluctuations in computing demand.
- The lock-in into proprietary subsystems, which impedes system evolution and integration that is possible under a modular architecture with appropriate application programming interfaces (API).

- Extending the data center operational systems with visualization, analytics and decision-support components that can be tailored to the needs of different stakeholders: business analysts, facilities engineers, IT administrators and cloud operators.

These challenges are key obstacles that deter the operators of small and medium data centers from experimenting with and adopting management techniques that capitalize on green energy. Having access to extensible, portable, and inter-operable solutions that address these challenges will facilitate cloud operators to experiment with numerous research outcomes[33,43-53] that propose alternative operational and pricing policies, to identify practices best suited to the specific configuration and usage of their facilities, and to develop, test, and deploy operational policies that optimize the efficiency and competitiveness of their services.

In this paper we present the design, implementation and deployment of, and early experiences with *ENEDI*, a modular cyber-physical system that we designed, developed, and deployed to address the challenges mentioned above. The ENEDI platform adopts a Microservice architecture[54,55] comprised of small, independent functional units deployed on Docker containers, which unify the management of a heterogeneous collection of IoT sensors and software agents that monitor the data center operation. ENEDI collects, integrates, and stores retrieved monitoring data, making them available for further analysis and visualization through the ENEDI dashboard. The dashboard is a Web-based application, which provides user-friendly configuration, visualization and decision-support functionality. ENEDI has been deployed and is in operation in the context of existing academic cloud data centers and adjoined solar energy production facilities, which were designed and installed to provision the data centers with renewable (solar) energy from photovoltaic (PV) panels .

The remainder of this paper is organized as follows: In Section 2, we present the steps taken to transform a conventional medium-scale academic data center into a "green" data center powered by solar energy; we analyse the technical, regulatory and financial considerations involved regarding the improvement of the energy efficiency of data center rooms and the procurement of a photovoltaic installation. In Section 3, we describe the physical and logical sensors installed to collect real-time information regarding the facility's operational conditions (electrical, environmental, computational) and status. In Section 4, we introduce the layered architecture of the ENEDI middleware platform, which comprises containerized microservices for retrieving data from the IoT sensors, and for integrated monitoring, configuration, discovery, and visualization of the data center's physical and IT resources. More details on implementation aspects of the platform are given in Section 5, and Section 6 presents early observations from an operational demonstrator of ENEDI, along with measurements and estimates on energy and cost savings and environmental impact. In Section 7, we give an overview of related work from the recent literature, and we conclude in Section 8 with main conclusions and discussion of future work.

2 | DATA CENTER TRANSFORMATION

2.1 | Overview

The transformation of a conventional ("brown") data center into one that is powered by Renewable Energy Sources needs to tackle a number of considerations related to the energy efficiency of the facility, the production and self-consumption of the energy from RES, the data-driven management of energy consumption and data center service provision, and the cost of the transformation effort (Figure 1).

The transformation process entails a number of necessary steps:

1. *Energy efficiency* measures must be examined, as they are of utmost importance to initially reduce a data center's energy consumption, especially with respect to its cooling needs, thus avoiding unnecessary investments in RES and/or other components, such as the expansion of cooling loads.[14,56] These measures can include the replacement of old cooling equipment or the installation of new equipment seeking to improve energy efficiency, the effective separation of hot and cold air streams with relevant materials, use of thermal insulation to reduce energy losses and so forth. This upgrade can provide a clearer insight for the RES system design which should come next, as the actual (reduced) data center energy needs will be taken into consideration, while ensuring energy conservation throughout its operation.
2. The design and implementation of a *renewable energy production* system that will produce on-site clean energy for self-consumption. Self-production of electricity via RES technology provides several technical advantages[47,57] and is encouraged by the rapid drop of RES prices[58-60] and various incentives established by regulators,[61-63] It assists in the reduction of a data center's exposure to the fluctuating retail electricity prices, thus enhancing the financial sustainability of the asset, and contributes to a reduction of CO_2 emissions.
3. Enabling *real-time monitoring* of all system components, data collection and in turn, data analysis. Availability of data allows *energy management and decision support* for optimized data center operation in terms of energy efficiency, economics and environmental conservation. Energy management and decision support will assist in enhanced operational flexibility, thus also increasing the financial sustainability of the asset.[15,44,64-66]

In the remainder of this section, we describe in more detail the steps that we took to transform a conventional, medium-sized academic data center in Cyprus, into a "green" facility that uses solar energy. We analyze the main considerations arising in the implementation of a collocated

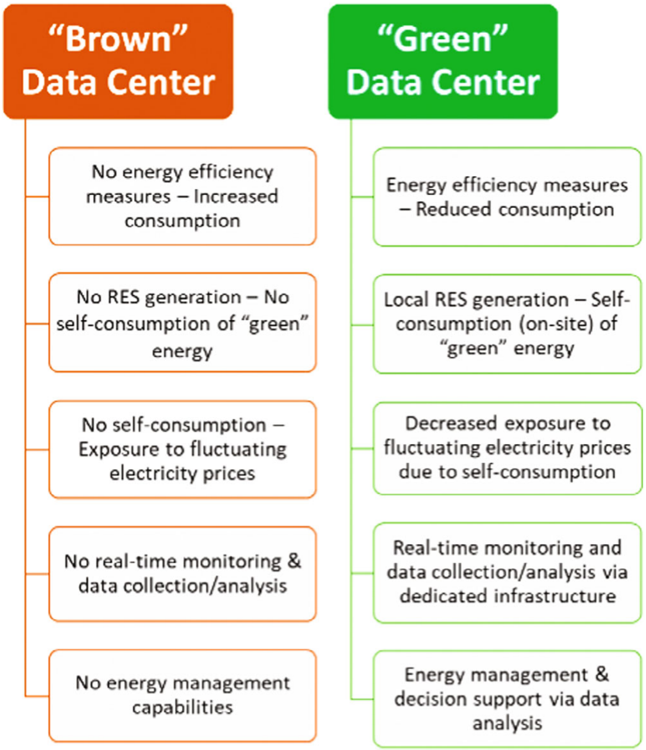

FIGURE 1 Comparison between a "brown" and a "green" data center.

PV facility, including regulatory and financial concerns pertinent to the European regulatory regime. We describe the steps taken to improve the energy efficiency of data center rooms (Section 2.2), and to implement a PV system on the building that hosts the data center. The procurement and installation the PV system was done in two stages:

1. *PV system design*, namely the determination of system capacity and geometric characteristics, technical specifications of equipment, system sizing and simulation with specialized software to estimate the monthly/annual energy production, Performance Ratio (PR), savings in CO_2 emissions, techno-economic study, electrical drawings, real estate plans and so forth (see Sections 2.3 and 2.5).
2. *Regulatory approval*, which entails the preparation and collection of necessary supporting documents and the submission of application forms to relevant local authorities (spatial planning bodies, energy regulator and distribution system operator) which issue the necessary licenses (see Section 2.4).

2.2 | Energy efficiency upgrade

The data-center's energy efficiency upgrade was primarily based on the collection of energy consumption data, as well as other electrical parameters of the site. For the spatial energy upgrade interventions, a new ventilator was installed for the improvement of cold air circulation and the abduction of hot air from the air conditioning systems, thus separating the cold from the hot air stream, increasing in turn the cooling system's efficiency. The data center hosts 14 computer racks, while the air conditioning is achieved by six direct expansion (DX type) close control units (CCUs). Prior to the ventilator installation, the supply and return of air to the space was free, without any air ducts.

The data center energy efficiency upgrade included the movement and/or rotation of the various air conditioning machines on-site, the rotation (or installation when absent) of cold air supply plenum boxes for providing the cold air and the configuration of the side covers of the machines so that the return is achieved from the opposite side of the supply.

Moreover, a special polycarbonate material was installed from the height of the existing racks to the ceiling to create an air barrier, so that the hot air does not mix with the cold one. The two narrow sides of the corridor remained open. Also, a metal air return duct was placed and securely fastened to the air conditioning units. The dimensions of the air ducts were determined considering a maximum allowable air speed of 10 m/s, as well as the balance of air flow between the units. Connection of the air ducts to the units with special anti-vibration material was also performed. The corners of the vents have turning vanes to facilitate the flow of air. The above aimed first to separate the cold from the hot air stream and, second, to surround the data-center's computing equipment by cold air, while ensuring that the waste hot air at the rear of the equipment keeps as low as possible.

FIGURE 2 Energy Efficiency Upgrade of data center room (equipment and positioning).

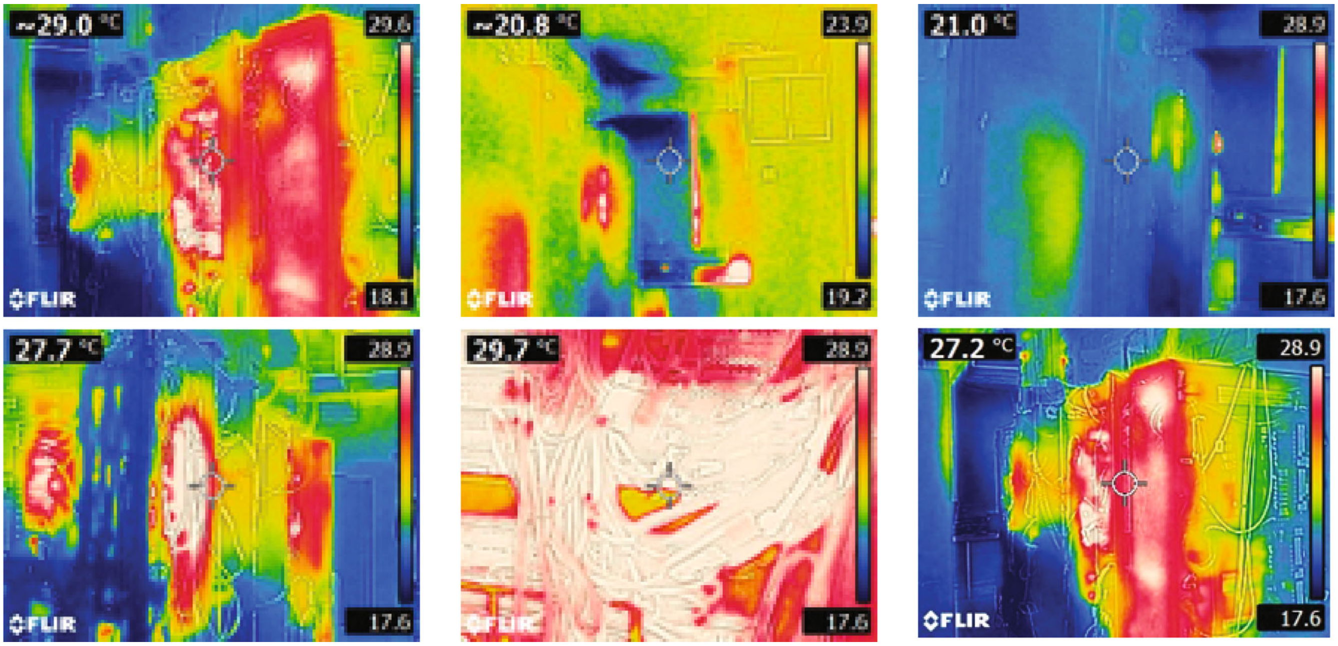

FIGURE 3 Thermal images taken after the completion of the energy efficiency upgrade. Clear separation of the hot air stream (red) from the cold one (blue), which increases the efficiency of the cooling system, reducing energy consumption.

Figure 2 illustrates the outcome of the energy efficiency upgrade. Figure 3 illustrates various thermal images taken in the data center, after the completion of the energy efficiency upgrade, detailed above. They clearly depict the separation of the hot (red colors) and cold air streams (blue colors) in the site, which increases the efficiency of the cooling system[56] and thus, reduces energy consumption.

2.3 | PV facility: Design considerations

The PV system design involves primarily the assessment of the environmental conditions of the installation location. Apart from the solar irradiance of the selected site and its geographical coordinates, other site particularities such as its orientation, its proximity to nearby shading obstacles, and the available area for use must be considered. During the design phase, compliance with the local technical requirements must also be ensured, specifically the system's inclination and orientation angle, its maximum height and distance from rooftop end (for rooftop PV installations), the allowable equipment positioning, as well as the maximum power capacity for grid integration, as required by the local authorities. Concerning the PV system sizing, the following parameters must be estimated with the use of a dedicated software tool: *Expected monthly and annual electricity generation* (kWh), *expected performance ratio* (%), *expected shading and other system losses* (%) and *expected savings in* CO_2 *emissions* (tnCO_2e).

Performance ratio (PR) is a widely used location-independent indicator serving as a quality factor, allowing the quantification of the good operation of a PV system. It is regarded as the relationship between the PV system's actual and theoretical energy yields (in kWh) for a specific period (usually a year), as seen in Equation 1. The theoretical energy yield is the expected energy produced according to the Standard Test Conditions (STC) as given by the manufacturer. A PR above 75% is generally considered as "very good."

$$PR = \frac{Actual\ PV\ production}{Theoretical\ PV\ production\ (based\ on\ STC)} \times 100\%, \tag{1}$$

where PR is the performance ratio indicator, actual PV production is the obtained energy yield of the PV system, and Theoretical PV production (based on STC) is the expected energy yield based on nameplate (STC) conditions as provided by the PV panels' manufacturer.

The CO_2 emissions avoidance indicator is a quantification index of the avoided CO_2 emissions due to a PV system's operation. It is calculated by multiplying the actual energy yield (in kWh) by the CO_2e grid factor of the country under study (in kgCO_2e/kWh), as seen in Equation 2.

$$CO_2\ emissions\ avoidance = Actual\ PV\ production \times CO_2\ grid\ factor, \tag{2}$$

where CO_2 emissions avoidance is the CO_2 savings indicator, Actual PV production is the obtained energy yield of the PV system, and CO_2 grid factor is an emissions index associated with each unit of electricity produced by a country's electricity system.

With regards to the case of residential PV systems, the objective is to establish the system with the maximum energy self-consumption.[67] In contrast, when sizing a PV system for commercial purpose, such as the case of a data center, the main aim is to implement the plant with the maximum electricity output, which will in turn result in the highest reduction of energy imports from the grid and thus the greatest decrease in electricity bills, with the lowest total investment cost though.

2.4 | PV facility: Regulatory issues and licensing

Grid-connected PV systems must comply with the local regulations and obtain the necessary license prior to the initiation of the work. Typically, these regulations are common among different countries, especially within the EU, despite various minor differences that can be observed. The necessary steps for licensing RES projects and particularly PV systems in Cyprus are summarized as follows:

1. Application to the Department of Urban Planning and Housing for their preliminary views.
2. Application to the Cyprus energy regulatory authority (CERA) for exemption from Construction permit requirement—for <5 MWp projects.
3. Preparation of Environmental Impact Assessment (EIA)—for >100 kWp projects.
4. Application to the Department of Urban Planning and Housing together with the EIA for the issuance of a Planning Permit.
5. Examination of the EIA by the Environment Committee and issuance of project's Environmental Terms.
6. Application for a Building Permit at the District Offices or Municipal Services that act as building authorities.
7. Application to the Distribution System Operator (DSO) or Transmission System Operator (TSO) for connection to the network (depending on the capacity of the project) and acceptance of the terms issued for the contract on payments and electricity purchases.
8. System installation, connection and application to CERA for Operating License.
9. Application to the local DSO or TSO for system inspection, testing and commissioning.

It can be seen from the above that the process of implementing, licensing and finally commissioning a PV system involves several local authorities and administrative procedures that can delay the plant's implementation. Upon the completion of the installation, inspection by the local DSO must be conducted to verify the compliance of the PV system with the local electricity regulations, which are commonly based on international standards. Specifically, for the case of Cyprus, the requirements of the German standard "Power Generating Plants Connected to the Low-voltage Grid (VDE-AR-N 4105)" must be ensured.[68] A summary of the Cyprus DSO requirements is listed in Table 1. These address the issues of over-/under-voltage and over/under-frequency, the protection of the equipment and the maintenance personnel after a possible loss of supply and its re-connection, as well as various power quality requirements.

2.5 | PV facility: Procurement and cost-benefit analysis

Following the analyses described above, the funding available to the project, and an investigation of available spaces compliant to building regulations, we embarked on the implementation of a PV facility at the rooftops of the building complex which hosts the data center. The configuration and placement of the PV panels was decided based on simulation results of solar exposure seeking to maximize production. We procured a 41.31 kWp rooftop PV system consisting of three different subsystems (17.01, 9.18, and 15.12 kWp, respectively). The ENEDI PV system comprises of 153 PV modules (JA Solar Cypress 270W Poly) of 16.51% efficiency, installed at an inclination angle of 17° to meet the local technical requirements, and three 3-phase PV inverters (FRONIUS SYMO), one for each PV subsystem. Inverters are responsible for converting the Direct Current (DC) power produced by the PV modules to Alternate Current (AC) power consumed by the building load or injected to the grid. The aforementioned subsystems' installation is shown in Figure 4.

TABLE 1 Electricity regulations for the connection of RES systems to the Cyprus Grid based on international standards.

Protection type	Regulation	Activation time
Under-voltage (U<)	0.8 × *Un* (i.e., 320 V)	1.7 s for RES systems with Low Voltage Fault Ride Through 200 ms for RES systems without Low Voltage Fault Ride Through
Over-voltage (U>)	1.10 × *Un* (i.e., 440 V) for RES systems <20 kWp 1.15 × *Un* (i.e., 460 V) for RES systems ≥ 20 kWp	200 ms
Under-frequency (f<)	47.0 Hz	200 ms
Over-frequency (f>)	51.5 Hz	200 ms
Protection against loss of main power supply (Loss of Mains – LoM): Voltage Vector Shifting islanding type and/or Active Anti Islanding according to standard DIN VDE 0126-1-1 (6-2006)	6 – 10° setup after operating control For the active protection of the Grid from islanding (active anti islanding protection), the Control and Protection circuit of the RES system performs voltage, frequency and complex resistance measurements, according to standard DIN VDE 0126-1-1 (6-2006)	200–1000 ms setup after operating system control
Re-connection time after the mains supply is restored	At least 180 s (>3 mins)—The increase of the active power that will be channeled to the Grid will not exceed the slope of 10% of the maximum active power for each minute	
DC current injection	<1% of the nominal phase current	<200 ms
Total Harmonic Distortion of output current (THDi)	<2% for RES systems connected to the Medium Voltage via transformer(s) <2.5% for RES systems connected to the Low Voltage	
Reduction of active power produced depending on the frequency, keeping the current value constant	When the grid frequency exceeds the limit of 50.2 Hz (fgrid ≥ 50.2 Hz) then the power produced by the system will be reduced by 4% for each 0.1 *Hz* increase in frequency (valid for 50.2 Hz ≤ fmains ≤ 51.5 *Hz*)	

FIGURE 4 Outdoor photovoltaic installation.

A cost-benefit analysis (CBA), which weighs the total system costs against its financial benefits, is important for assessing the economic feasibility of such investments prior to their implementation. The total system cost includes the initial cost, which comprises of the installation cost (in relation to the system size) and the connection cost, and the annual operational cost, which further includes the operation and maintenance cost, the insurance cost and other size-dependent and/or size-independent costs. The financial evaluation of the ENEDI PV system is based on the calculation of three fundamental financial indicators, namely the *net present value* (NPV), *Payback Period* (PP) and *Levelised Cost of Electricity* (LCOE), which are widely exploited for the assessment of such investments.[59] Specifically, the NPV (in €) is the deviation between the present value of cash inflows and outflows over the investment's lifetime (t) considering a specific discount rate (d), as seen in Equation 3. In general, a positive NPV depicts a profitable investment.

$$NPV(t) = \sum_{n=1}^{t} \frac{Cash\ Flow}{(1+d)^n} - \left(Initial\ Cost + \sum_{n=1}^{t} \frac{Operational\ cost}{(1+d)^n} \right), \tag{3}$$

where NPV is the Net Present Value indicator, t is the system's lifetime, n is the reference year, d is the discount rate, Cash Flow is the amount of money equivalent generated by the system, Initial Cost is the system's total investment cost, and Operational Cost is the system's yearly cost of operation.

The PP (in years) is the period needed to break-even the investment, as seen in Equation (4). Generally, a low PP reveals a more attractive investment to the investor.

$$NPV(PP) = 0 \Rightarrow \sum_{n=1}^{PP} \frac{Cash\ Flow}{(1+d)^n} = \left(Initial\ Cost + \sum_{n=1}^{PP} \frac{Operational\ cost}{(1+d)^n} \right), \tag{4}$$

where NPV is the Net Present Value indicator, PP is the Payback Period indicator, n is the reference year, d is the discount rate, Cash Flow is the amount of money equivalent generated by the system, Initial Cost is the system's total investment cost, and Operational Cost is the system's yearly cost of operation.

Finally, the LCOE (in €/kWh) is the ratio of the present value of total costs to the total energy produced,[59] as seen in Equation (5).

$$LCOE(t) = \frac{Initial\ Cost + \sum_{n=1}^{t} \frac{Operational\ cost}{(1+d)^n}}{\sum_{n=1}^{t} \frac{PV\ Production}{(1+d)^n}}, \tag{5}$$

TABLE 2 ENEDI PV system's estimated financial indicators for a 20-year operation.

Financial parameter	Estimated value
Net Present Value (NPV)	€92,914.76
Payback Period (PP)	4 years
Levelised Cost of Electricity (LCOE)	0.055 €/kWh

where LCOE is the Levelised Cost of Electricity indicator, t is the system's lifetime, n is the reference year, d is the discount rate, Initial Cost is the system's total investment cost, Operational Cost is the system's yearly cost of operation, and PV Production is the obtained energy yield of the PV system.

The ENEDI PV system is expected to generate at least 66.52 MWh of "green" electricity during its first year of operation. This results in €13,171.82 money savings, considering the average price per unit of electricity of the University of Cyprus (UCY), that is, 0.198 €/kWh incl. 19% VAT, as calculated according to 2019 electricity bills. Accounting the initial installation cost of €52,734.85 incl. 19% VAT (which complies with the benchmark value for PV installations in Europe[60]) and assuming a 20-year system operation given the system's warranties, the analysis results are summarized in Table 2. It must be noted that the 5.5 c€/kWh ENEDI PV system's LCOE is very close to values observed for other PV investments in Cyprus, specifically 5–7 c€/kWh, which are among the lowest in Europe,[60] thus confirming the attractiveness of such investments in Cyprus and mostly, the economic viability of the ENEDI PV system.

3 | IOT SENSOR INFRASTRUCTURE

In this Section we describe the sensing infrastructure put in place to collect real-time data regarding the facility's operational conditions and status.

3.1 | PV facility monitoring: Outdoor sensing infrastructure

The PV system comes with embedded sensors, which capture its output and operational conditions in real time (Table 3, left column). Additional IoT sensors are installed with the PV system to monitor the environmental conditions that affect PV production (Table 3, right column).

We experimented with three alternative configurations for connecting the outdoor PV and environmental sensors of Table 3 to the local network and extracting their metrics (see Figure 5): (i) a special *data logger* device (Campbell Scientific - CR1000), equipped with its own HTTP Module; (ii) a low-cost configuration of bespoke *Arduino* and *Raspberry Pi* micro-controllers, and (iii) connection to the PV inverters and extraction of metrics through their embedded HTTP servers. We decided to use the PV inverters to retrieve the metric streams collected by the PV-embedded sensors, as this solution provided better reliability without extra cost.

3.2 | Data center room monitoring: Indoor IoT sensing infrastructure

To collect data from systems inside the data center, we installed an indoor sensing infrastructure that monitors electricity consumption and relevant environmental conditions near the data center's racks. This infrastructure consists of electricity distribution units with integrated smart electricity meters and appropriate software for data collection, analysis and distribution, as well as dedicated environmental sensors.

A total of 14 *Smart Power Distribution Units (SPDUs)* were installed at the data center's server racks. The SPDU is a stand-alone power distribution device that provides the ability to remotely monitor connected loads in real time, as well as user-defined alerts for possible circuit overloads. The

TABLE 3 Parameters collected from the outdoor sensing infrastructure.

PV Sensors	Environmental sensors
DC & AC power (W)	Ambient Temp. above the PV modules (°C)
DC & AC current (A)	Ambient Temp. below the PV modules (°C)
DC & AC voltage (V)	Ambient Humidity in %
Energy produced (kWh)	Wind Speed in km/h
Energy injected to the Grid (kWh)	Global Solar Irradiation in kWh/m^2

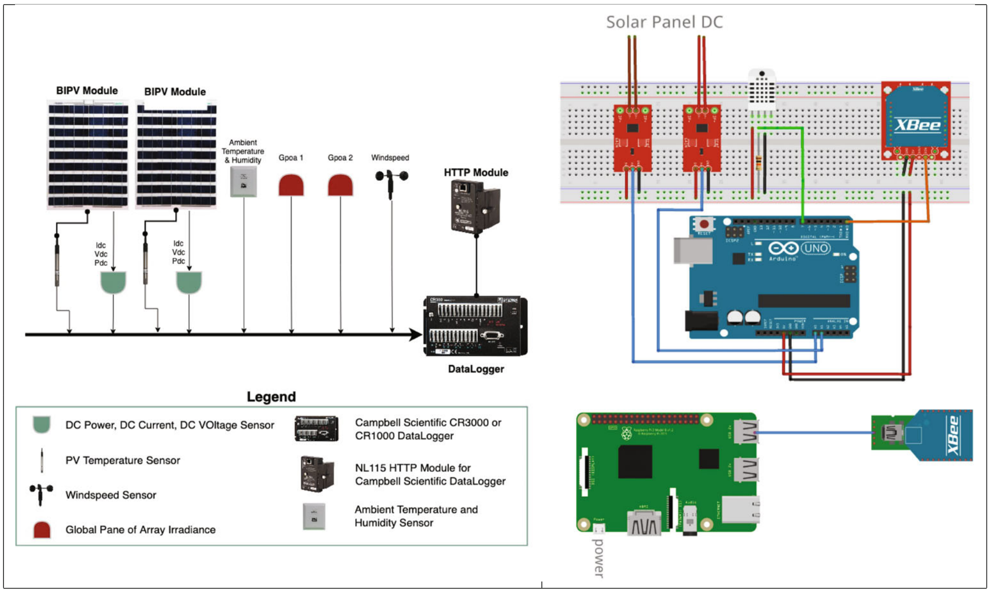

FIGURE 5 Sensors for outdoor PV monitoring.

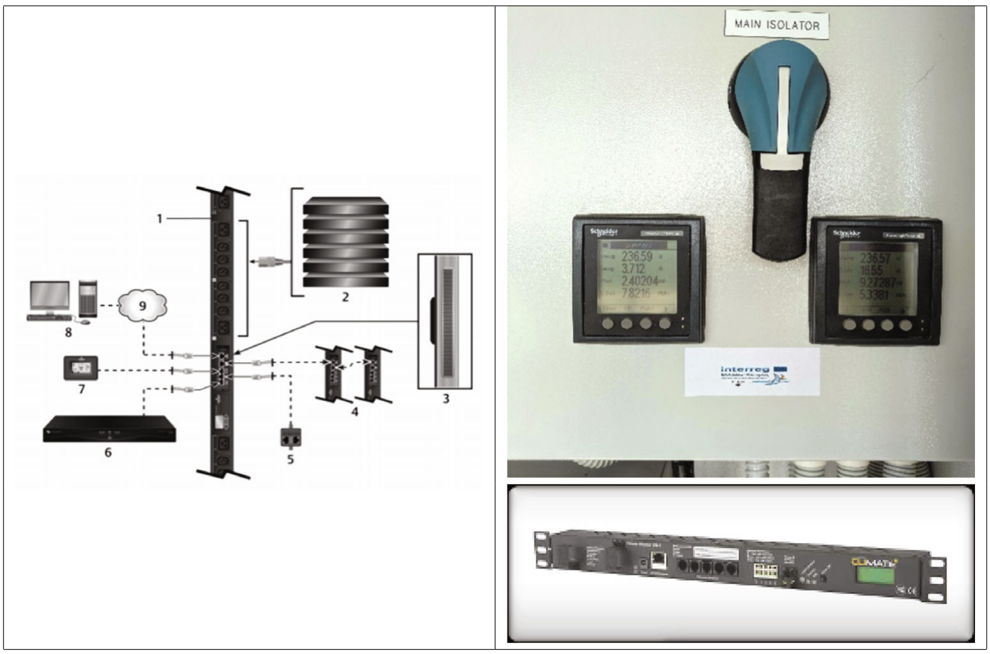

FIGURE 6 Indoor electricity and climate sensors.

SPDU can achieve power consumption control with integrated and flexible monitoring capabilities by continuously measuring and recording power consumption (1-phase or 3-phase) via a secure Web or Simple Network Management Protocol (SNMP) interface. We selected the Liebert Power Distribution MPH2—Input measurement, output measurement, output switching (Type R), as shown in Figure 6 (left) and Table 4 . This SPDU has an input rating of 230 V, 16/32 A, 24 output sockets, size of $56 \times 60 \times 1736$ mm, and it is able to measure the following parameters: Current (A), Voltage (V), Real Power (W), Apparent Power (VA), Consumption (kWh), Power Factor and Frequency (Hz), with an accuracy of $\pm$ 1%. The selected SPDU offers also the possibility of supporting temperature and/or humidity sensors for relevant measurements at the rack level, at both the cold air intake point and the hot air discharge point. Therefore, we also installed 14 joint temperature/humidity sensors on our racks. Specifically, two CLIMATE CM-2 sensors were installed along with the SPDUs for every two server racks, one on the front side (cold zone) and the other on the rear side (hot zone) of the rack (see Figure 6 (bottom, right)). Finally, to monitor the electricity consumption of the cooling systems, we installed smart electricity monitors with the data-center's air conditioning units (see Figure 6 (top, right)).

TABLE 4 Liebert power distribution MPH2 SPDU description [source: vertiv.com].

Item	Description
1	Vertical rack PDU
2	Connected equipment
3	Case ventilation, both sides (optional)
4	Rack PDU array
5	Sensors—Integrated and modular
6	Serial appliance
7	RPC basic display module (BDM)
8	Monitoring station
9	Network connection (10 MB/100 MB/1 GB)

In summary, the ENEDI Indoor IoT Sensing Infrastructure consists of the following subsystems:

- Subsystem 1: Remote monitoring and recording of the electricity consumption of the data center's server racks, as well as their environmental conditions (temperature and humidity).
- Subsystem 2: Remote monitoring and recording of the electricity consumption of the data center's cooling system (air conditioners).
- Subsystem 3: Remote monitoring and recording the data center's (indoor) environmental conditions (temperature and humidity).

All measurements recorded by the aforementioned metering devices are collected by ENEDI's Monitoring System through the local IP network network (see Figure 8 in the next section).

3.3 | Computing cluster virtualization and monitoring

The data-center's clusters comprise computing and storage servers organized in racks and connected via a high-speed LAN. The installation and configuration of the physical (bare-metal) servers' *host* operating system (Debian-based Linux) and of essential cluster services (DNS, DHCP, LDAP, etc) are performed automatically through the *Metal-a-a-Service* (MaaS) toolset.[69]

We adopt virtualization as the key enabler for offering cloud services according to the Infrastructure-as-a-Service (IaaS) delivery model and install the Kernel-based Virtual Machine (KVM) hypervisor[70] with the Linux OS of our cluster bare-metal servers. KVM is a Type-I ("bare-metal") hypervisor built into the Linux kernel that can host virtual machines (VMs) with different "guest" operating systems. Virtualization brings a number of significant advantages, such as abstraction, server consolidation, performance isolation, and easier resource provisioning.[5,71,72] To take advantage of virtualization and facilitate the management of VMs, we deploy *Openstack*,[73,74] the open-source IaaS middleware management system.

Applications deployed on ENEDI's infrastructure can be packaged and launched as Docker Containers[75] deployed on VMs, as this provides significant benefits: simplified application deployment, good isolation with low overhead and short booting times.[76-78] Docker is deployed inside Linux VMs, which provide support for the Docker Engine.[75]

To make the performance and energy consumption of physical servers obtainable, we installed on their host OS the *Ipmitool*[79] and a *Netdata agent*.[80] *Ipmitool* is an open-source utility supporting the *Intelligent Platform Management Interface* (IPMI) open standard,[81] which enables remote server monitoring and control. We use the Ipmitool to retrieve information about the status of physical servers, including system temperatures and power consumption. *Netdata*[80] is a state-of-the-art, open-source, distributed platform for monitoring systems and applications in real time. The Netdata agent is a lightweight metrics collector, which can collect thousands of infrastructure metrics from a host OS, without disrupting its operation or hurting its performance. The agent exposes the performance metrics it collects through the Netdata API.[80] ENEDI's users can package their applications in ENEDI-ready Docker containers with preinstalled Netdata agents, which offer application-level monitoring. The configuration of ENEDI's data center clusters is summarized in Figure 7. More information on the monitoring system is given in Sections 4 and 5.

4 | PLATFORM ARCHITECTURE

4.1 | Overview

To maximize the environmental and financial benefits derived from a PV installation that provides electricity to an associated "green" data center, the data-center's operations and cloud-service provisioning need to take into account factors that affect both the short- and long-term trends

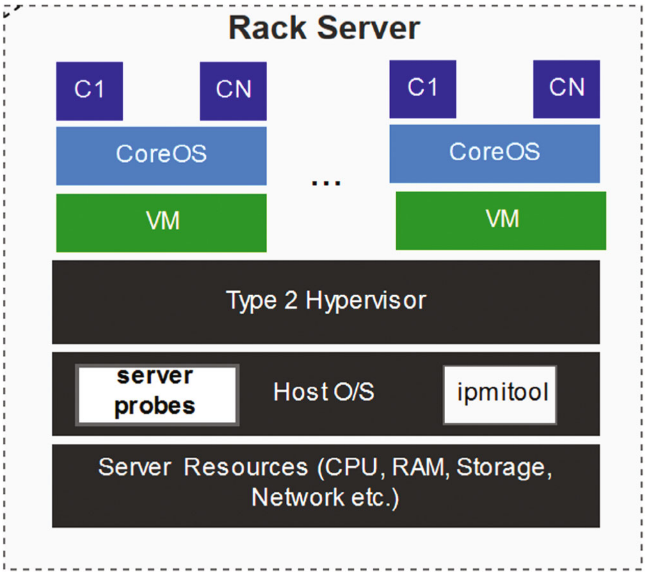

FIGURE 7 ENEDI server virtualization layers.

TABLE 5 Mapping the ENEDI architecture to the IoT reference architecture modules.[84]

IoT reference modules	Infrastructure layer	Platform layer	Dashboard
Interoperability	✓	✓	
Persistence		✓	✓
Analytics		✓	✓
Context			✓
Resources	✓	✓	✓
Events	✓	✓	✓
GUI			✓
Security	✓	✓	✓

of solar-energy production and the electricity consumption of the data-center's equipment. To make this possible, we designed and implemented ENEDI, an integrated modular software platform for managing the various IoT sensors described earlier. ENEDI retrieves sensor data on a continuous basis, and makes near-real-time and historical information available to data center and cloud-service operators, end-users, and other subsystems for further analysis and decision support. ENEDI's architecture aims at separating functionalities across different layers, hiding the heterogeneity of hardware, software, and communication protocols behind simple abstractions and interfaces.

ENEDI is designed as a three-tier system comprising an IoT *Sensor Layer*, a *Platform Layer* and a *Dashboard*. This architecture is in-line with reference architectures of edge computing systems[82,83] and IoT Middleware.[84] In particular, the functionalities encapsulated in ENEDI's layers cover the main requirements considered important in the IoT Middleware reference model of Reference 84, summarized in Table 5. The three layers of the architecture are described in more detail below.

4.2 | Sensor layer

The main purpose of this layer is to facilitate the integration and configuration of sensors described in Section 3 in order to achieve the seamless retrieval of measurements (metrics) representing the status, operation, and environmental conditions of the "green" data-center's facilities. Sensors are distinguished into two categories: (A) *Physical (hardware) sensors*. These are embedded devices, which capture: (i) The solar energy production in associated PV facilities; (ii) Environmental conditions inside and outside the data center (temperature, humidity, solar radiation, wind speed, etc.); (iii) The energy consumed by data center facilities: computing servers, server racks, clusters, and server room air-conditioning systems. The physical

TABLE 6 Groupings of physical (hardware) sensors.

Monitored Infrastructure	Edge device	Protocol used	Location of Edge device
PV panels	DataLogger or Inverter	HTTP or ModBus	Rooftops
Data center conditions	ClimateMonitor-2	HTTP or SNMP	Cluster rooms
Power consumption	Smart PDUs	HTTP or SNMP	Server racks

TABLE 7 Physical and logical sensors of ENEDI.

Sensor	Metric unit	Type	Frequency
Physical sensors (Location: PV installation)			
Pyranometer	GPOA	Env.	15 s
DC voltage	Vdc	Perf.	15 s
DC current	Idc	Perf.	15 s
Ambient temp. & humidity	∘C, %	Env	15 s
PV temperature	∘C	Env.	15 s
Wind speed	km/h	Env.	15 s
Physical Sensors (Location: Data center Rooms)			
Temperature	∘C	Env.	15 s
Humidity	∘C	Env.	15 s
Dew point	∘C	Env.	15 s
Light	%	Env.	15 s
Sound Level	dB	Env.	15 s
Water Sensor	%	Env.	15 s
Current	Amp	Perf.	ms
Power Factor	[0–1]	Perf.	ms
Logical sensors (Location: Computing servers)			
IPMI: temp., fan, volt.	∘C, rpm, V	Env.	ms
Netdata	Various metrics for server	Perf.	ms
Probes	cluster, application resources		
Custom Prometheus	Various metrics for VM,	Perf.	ms
collectors	Container & Application resources		

sensors are grouped according to the facility infrastructure they monitor (PV installation, cluster rooms, etc.). The sensors of each such group are typically connected to an *edge device* through wireless (Wi-Fi, Bluetooth) or wire-line connection. This device acts as a gateway, providing connection to the local network, retrieving sensor metrics using the appropriate protocol (HTTP, SNMP, Modbus), and feeding the collected metrics to the platform's monitoring system. Table 6 summarizes the grouping of physical sensors of the ENEDI infrastructure. (B) *Logical (software) sensors*. These are implemented as software agents that retrieve metrics from the operating systems of the data-center's physical or virtual machines, the run-time environment of applications running on the data-center's computing resources, or third-party APIs available through the Internet. Logical sensors collect data about (i) the utilization of data-center servers and networks; (ii) the profiles of workloads accepted for execution which contribute to the data-center's power consumption; (iii) operational expenses and income from paying customers; (iv) weather conditions and weather prediction; (v) pricing information about the electricity purchased from the electricity grid and so forth. Metrics retrieved are fed into the platform's monitoring system.

Tables 6 and 7 summarize the physical and logical sensors of ENEDI and Figure 8 presents an abstract diagram of the various components that comprise the sensor layer of our installation.

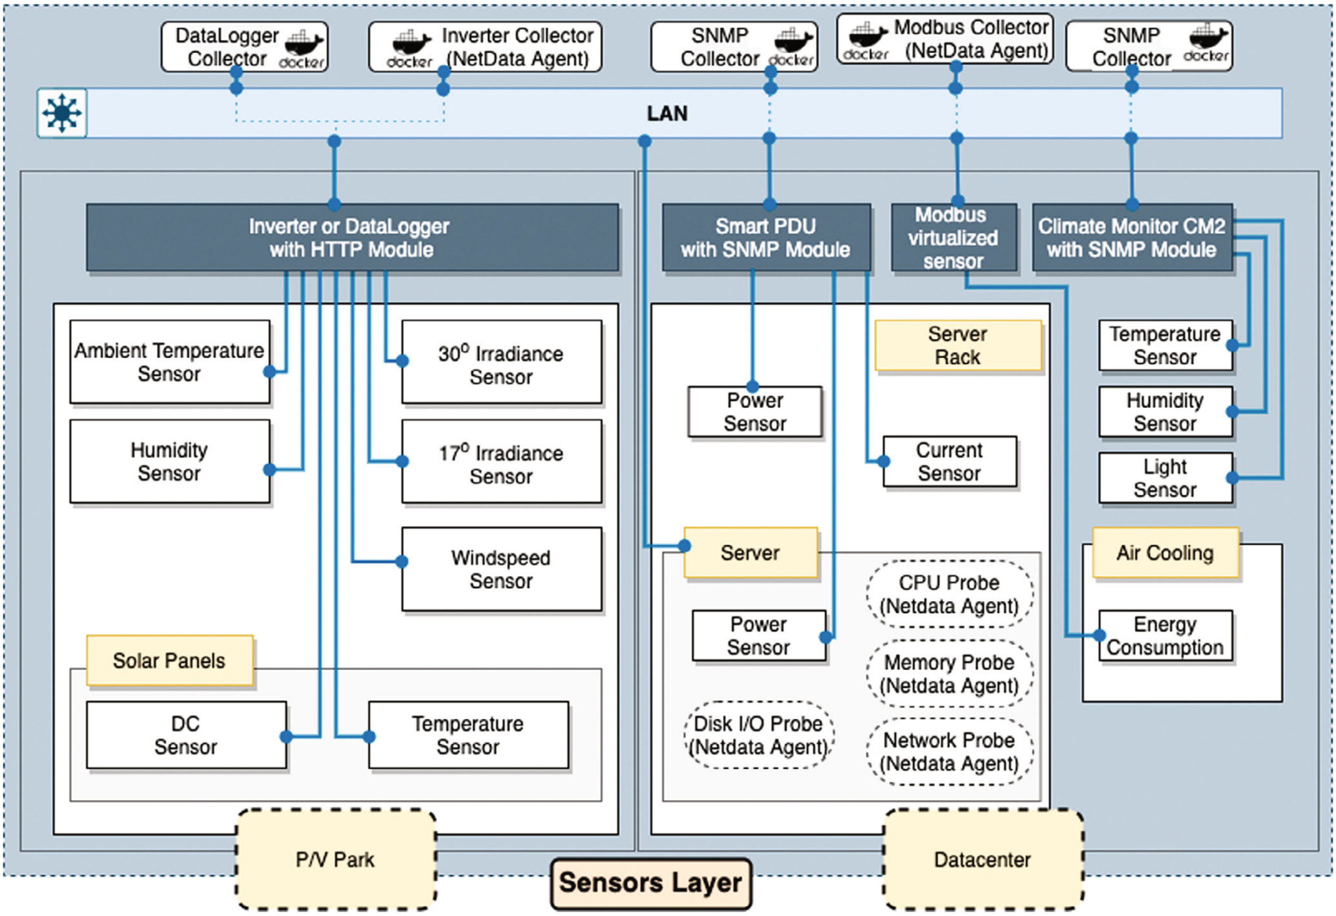

FIGURE 8 Sensor layer: Physical sensors are represented as white rectangles, logical Sensors as ovals, and edge devices as grey rectangles.

As can be seen from Table 7 and Figure 8 the hardware and software sensors employed have diverse characteristics and speak different protocols. Therefore, the goal of the Sensor Layer is to transform these diverse systems into an abstract sensing infrastructure that is easily configurable, manageable, and which supports the simple abstraction of feeding sensor data to the ENEDI Platform through common interfaces. Each sensor that is deployed at the data center facility is paired with a *Collector*, namely a small program configured to interact with a physical or logical sensor, retrieve its metrics, and make them available to the ENEDI platform's monitoring service with appropriate tagging for proper identification, further dissemination, storage, and analysis. We use containerization[78] to implement this layer and each Collector program is configured and deployed to run inside its own Docker container[75] (see Figure 8).

4.3 | ENEDI platform layer

This layer comprises services for the management of the data center and its cloud operations. It is responsible for the retrieval of data from the Sensor Layer and from third-party services, and the integration, storage, and provision thereof to operational subsystems and top-layer modules. For the design and implementation of the ENEDI platform, we adopted the microservices architectural paradigm, which has become quite popular for the development of cloud-native applications[54,55,85] and cloud infrastructure services.[86,87] According to this paradigm, applications are decomposed into small functional units, each with a discrete business capability, and are delivered as services which run as independent processes and communicate with each other and the outside world through well-defined, platform-agnostic APIs and lightweight communication protocols like HTTP.

Microservices is certainly not a new approach as it drives ideas from Service-Oriented Architectures.[88] However, if placed in the context of lightweight virtualization technologies[86] and the need to integrate application development with application deployment and management through modern DevOps techniques,[55,89,90] it offers many advantages including:

- The capability to easily integrate and reuse service components developed in different programming languages and running on different run-time environments.[86]

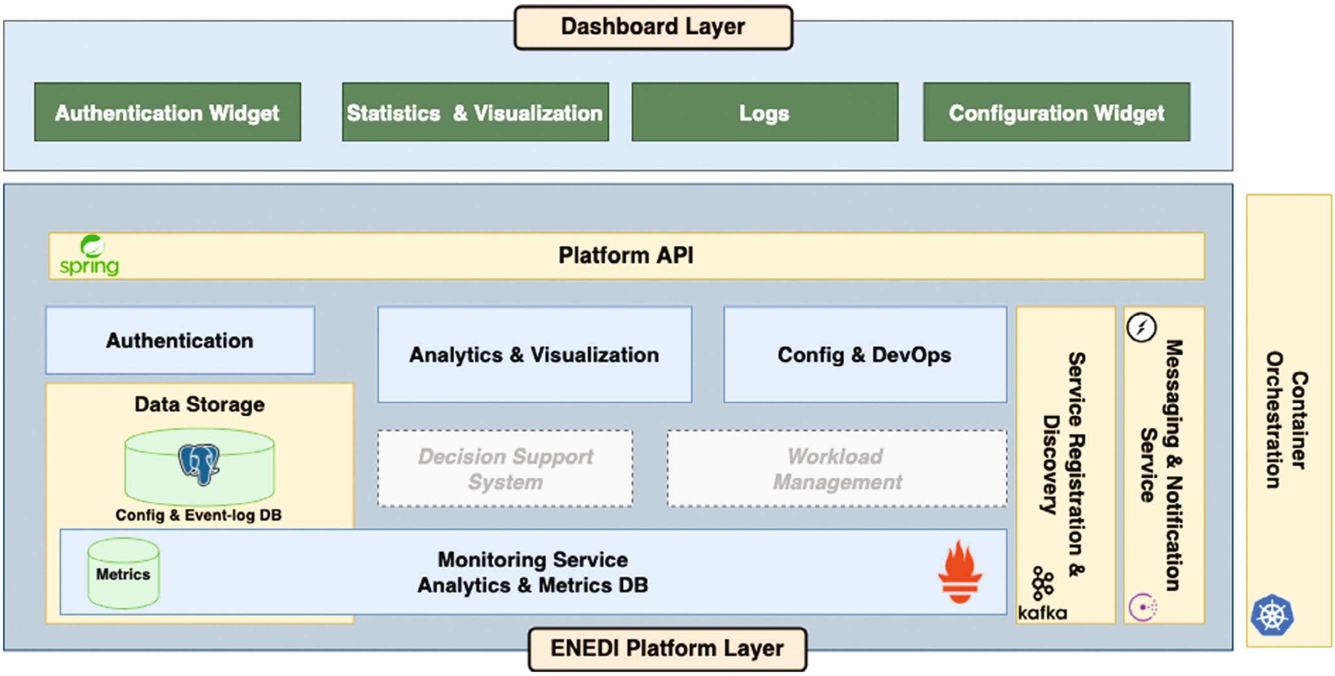

FIGURE 9 ENEDI layers: Platform (middle) and Dashboard (upper).

- Improved efficiency and agility since containerized microservices run without the overhead of a guest OS, have shorter start-up times, are less memory-demanding and easier to migrate.[86]
- Enhanced management efficiency as microservices can be independently deployed, monitored, upgraded, and scaled, enabling the uptake of continuous integration and delivery (CI/CD) and automated performance and failure management techniques.[55,88,89]
- Easier evolution as new functionality modules can be deployed without having to undertake changes in running microservices.
- The capability to select and use alternative software components, such as monitoring library, database, or messaging middleware, according to the preferences of each installation.
- Easier alignment between business goals and IT operations can be developed and managed independently by specialized teams of DevOps engineers and business analysts.

Infrastructure microservices provide implementation and support for *service discovery* and *orchestration*, *messaging* and *notifications*, and *persistent storage* for configuration files, logs, and streams. Higher-level microservices support functionalities, which are currently offered through the ENEDI Dashboard, namely *User Authentication*, *Data center Monitoring*, *Analytics and Visualization*, *Configuration Management and DevOps*, or are planned for a future release (*Decision Support*, *Workload Management*). Last, but not least, the ENEDI platform comprises an *API Gateway*, which makes its services accessible through RESTful APIs. A high-level overview of the ENEDI platform's architecture is shown in Figure 9. More details on the implementation of the ENEDI platform are given in Section 5.

4.4 | Conceptual data model

The dashboard and the monitoring system share a common conceptual data model, which defines the logical organization and the naming conventions of the metric streams retrieved from the Sensor Layer. This model reflects the architecture of the data center and the hierarchical grouping of its IT, power-production, and cooling equipment. The naming defined by the model is used to annotate collected metric streams produced by the sensors, retrieved from and disseminated by the monitoring system.

In practice, the model is "populated" with properly organized and annotated data retrieved by the Dashboard application through REST calls to the ENEDI platform's API server and respective queries to its monitoring system. Retrieved data represent the configuration of the infrastructure (clusters, PV components, connected sensors) and the streams of readings from sensors, properly grouped according to the data center's architecture. The adoption of this conceptual data model by geo-distributed data centers induced in the ENEDI platform, enables the use of the ENEDI Dashboard as a shared data platform that integrates monitoring and configuration data retrieved from collaborating data centers.

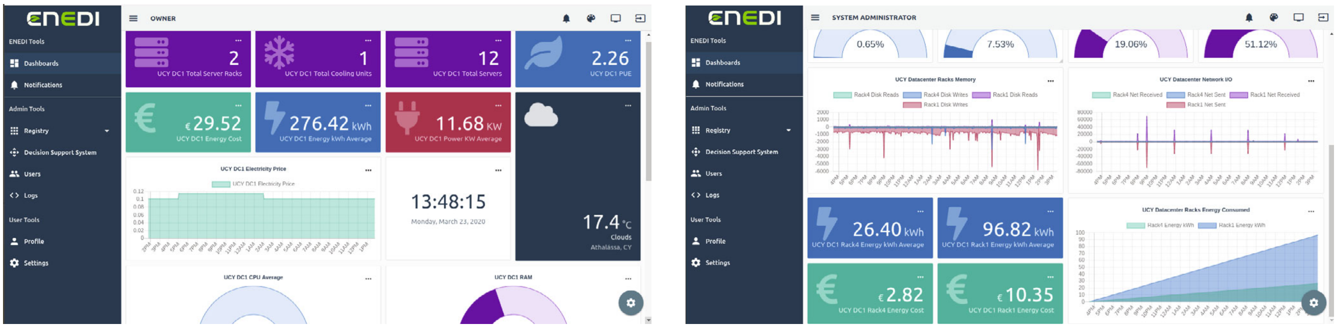

FIGURE 10 Examples of ENEDI dashboard of business analyst (left) & system administrator (right).

4.5 | Dashboard layer

At the upper layer of the ENEDI architecture lies a modular Web-based *Dashboard*, which provides the following functionalities: (i) user authentication and management, (ii) data center monitoring Graphical User Interface (GUI), (iii) data visualization and analytics, and (vi) configuration of sensors, resources and services of the platform. The dashboard's GUI comprises a number of widgets, which can be configured and adapted by its user. The organization of the GUI reflects the logical structure of a data center or a federation of data centers, with each data center comprising one or more clusters, cooling equipment, and a solar energy production facility, each cluster comprising one or more racks, each rack comprising a number of servers, and so on. The dashboard provides preset configurations that are customized to the needs and privileges of three groups of data center stakeholders: (i) *Data center Business Analysts/Owners*: these are users who need to have access to consolidated information regarding the overall use of data center IT resources, power consumption, cost of electricity, solar energy production, energy self-consumption ratio, PUE, CO_2 emissions, total operational expenses of the data center and the evolution thereof with time, Return-on-Investment of renewable energy, environmental impact of using renewable energy and so forth (see Figure 10, left). *System Administrators*: these users need access to more detailed, technical information about the status of the facility (temperature, humidity, power consumption, CO_2 emissions) and its computing and networking resources (servers, server-racks), the cloud services in operation, workloads deployed on the cloud infrastructure and their energy profiles and so forth. (see Figure 10, right) *End Users*: these are users of the data center's cloud services who need to have access to information about the computing resources consumed, money spent (where applicable) and the environmental and energy profile of their workloads.

The dashboard is structured as a set of modular front-end components, in order to facilitate integration with third-party libraries and APIs. Each integration and User Interface (UI) element can be incorporated as an individual extensible component. The dashboard are delivered as a Web Single Page Application (SPA)[91] developed with the Angular[92] framework and deployed on a dockerized Nginx Web server.[93] The dashboard uses the ENEDI platform's notification service (Websockets) to dynamically update its visualizations when dataset updates are detected on the back-end.

The dashboard offers also a suite of Web-based forms that facilitate the configuration of the ENEDI platform, sensors, and visualization widgets. These forms allow a user to input the necessary configuration information into the system to specify alternative dashboard visualizations and register new components, namely data centers, clusters, computing servers, edge devices, sensors or renewable-energy sources. The dashboard translates the configuration information inserted into appropriate configuration instructions expressed as Docker commands, which can be launched to the ENEDI platform's environment in order to produce the required configuration actions. Figure 11 presents representative screenshots from the dashboards configuration capabilities, specifically when an operator seeks to induce a new data center on the ENEDI platform.

5 | IMPLEMENTATION AND DEPLOYMENT ASPECTS

In this section, we provide more details on the implementation of ENEDI's platform.

5.1 | Infrastructure services

A number of libraries are incorporated in the ENEDI Platform to provide services required to deploy and/or implement core functionalities, such as service discovery, configuration, messaging and persistent storage (see Figure 9).

As mentioned earlier, each ENEDI microservice runs inside a separate Docker container,[75] which is deployed on a virtual machine of the data-center's *Openstack*[73,74] virtualization platform. To facilitate the orchestration and management of these containers, we employ

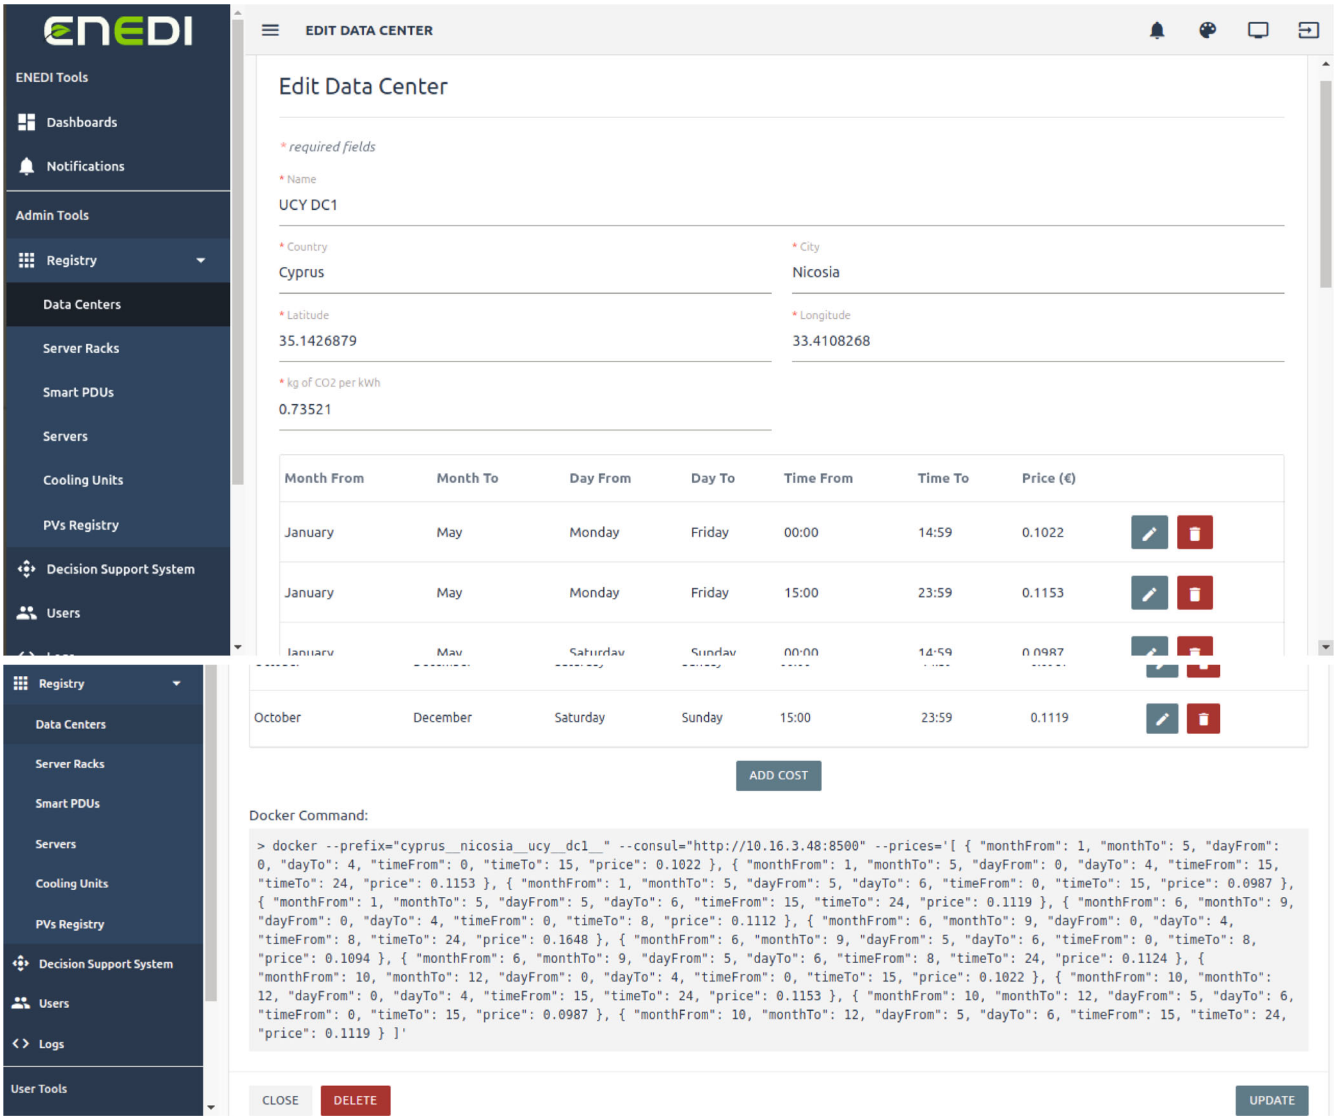

FIGURE 11 ENEDI Dashboard configuration functionality: adding a new data center to the system.

the *Kubernetes* orchestration platform[76,94] and deploy it on an Openstack VM. We manage Kubernetes using Rancher,[95] a tool offering GUI-based access to Kubernetes functionality for virtual cluster configuration, deployment and management of Pods where ENEDI's Docker containers run.

ENEDI microservices need to be configured properly so they can locate each other and communicate seamlessly among them. Moreover, new sensing, energy and computing resources should be integrated to the platform without having the IT operator to manually stop and reconfigure running services. To this end, we employ a Discovery and Configuration system based on Hashicorp Consul,[96] which supports the dynamic registration, configuration and discovery of ENEDI microservices. When a new microservice is launched inside its container, it registers with Consul and receives configuration information. Consul acts as a service registry, enabling microservices to seek and discover their counterparts before engaging in information exchanges, and performing health checks of registered services. If there is a configuration change at run-time, the Consul registry is updated and notifies appropriate microservices dynamically.

Changes in the status of the data-center's resources have to be propagated to the ENEDI Dashboard in order to update accordingly its UI. Propagation of changes is undertaken by the *Notification Service*, which uses WebSockets[97,98] to connect with an open socket of the Dashboard and push changes there in near real-time. The Notification Service comprises also a Message Queue implemented on top of *Apache Kafka*,[99] on which platform components (producers) can push messages on a specific topic so that other components (consumers) receive, process, and store the messages. Finally, the ENEDI Platform comprises services for the persistent storage of configuration files and logs on a Postgres relational database,[99] of raw metrics coming from the Monitoring System.

5.2 | Monitoring system

The *Monitoring System* is one of the key components of the ENEDI Architecture. The goal of the monitoring system is to provide for the continuous retrieval, integration, dissemination, and filtering of metric streams produced from physical and logical sensors.[100-102] Several monitoring systems are available for use in cluster and cloud computing environments, both open-source and proprietary, which offers the following core functionalities:[100,103,104] (i) *Configuration*: this deals with setting up the monitoring infrastructure and connecting together its components (sensors, probes, agents) that produce, collect, and disseminate the metric streams sought. (ii) *Collection*: this addresses the actual retrieval of sensor data, in line with the properties of sensors as defined during the configuration process, namely the access method (push vs. pull), the structure and naming of metrics (key-value pairs or groups of metrics), the sampling period and so forth. (iii) *Dissemination*: this refers to the transfer of sensor data from the edge of the system (e.g., data-center servers, data centers and PV Collectors as shown in Figure 8) to a central Monitoring Service. This is often done using REST API calls in combination with a Message Queue that follows the pub/sub protocol (Probes/Collectors are publishers and the Monitoring Service is the subscriber). (iv) *Storage*: the Monitoring Service stores the collected sensor data as metrics streams in a time-series database for historical and analysis purposes.

The monitoring system of ENEDI is deployed and managed as a containerized microservice, which is part of the ENEDI Platform (see Figure 9). We develop ENEDI's monitoring system on top of Prometheus,[103] an open-source, industry-leading toolkit. Prometheus provides a multi-dimensional data model for time-series of key/value pairs identified by metric name, a flexible query language to leverage this dimensionality, no reliance on a particular distributed storage solution (a Prometheus server can store retrieved metric streams), and autonomous single server monitoring nodes. Prometheus provides also libraries for developing *instrumentation probes* in many popular programming languages (Python, Go, Java, Ruby, C, C++, etc.). This facilitates the development and integration of Prometheus-compatible instrumentation probes with different sensor devices (these are referred to as *ENEDI Collectors* in Section 4.2). *Netdata*[80] supports integration with Prometheus thereby enabling the seamless retrieval of sensor data collected by Netdata agents that run on various edge devices, without further programming effort. The monitoring data produced by Prometheus' instrumentation probes are pulled by Prometheus servers using a pull model over HTTP, although a push model is also supported through intermediary gateways. The discovery of instrumentation probes is done via service discovery. In the ENEDI implementation, Collectors register their presence with Consul,[96] at configuration time, and the Prometheus server identifies their presence by querying Consul. Table 8 presents the metric streams collected by Prometheus in ENEDI's monitoring system.

Finally, Prometheus supports the hierarchical federation of Prometheus servers in a tree-like topology where higher-level Prometheus servers collect aggregated time-series data from a larger number of subordinate servers. This feature allows Prometheus to easily scale to environments with multiple geo-distributed data centers, which share a common conceptual data model, and thousands of sensors. In particular, the hierarchical federation mode of Prometheus acts as a convenient representation of geographically distributed ENEDI enabled data centers. Each participating data center has its local subordinate Prometheus instance that collects all related metrics produced by the agents and sensors deployed on site. Then, the central Prometheus, through the federation mode, scraps monitoring data upon configurable time intervals, thus making all monitoring data from all sites centrally available to the ENEDI Platform. This approach offers also the benefit of on-the-fly registration and de-registration of ENEDI data centers dynamically without prior knowledge of the number or location of participating data centers. In addition, each data center site's Prometheus prefixes its metrics with a prefix unique to the specific geo-located data center thus making it easy to query metrics from specific locations

5.3 | API gateway

The API Gateway is an important component of a microservices-based system like ENEDI. The purpose of the Gateway is to shield "back-end" microservices from the outside world and provide access to their capabilities through a carefully designed RESTful API.[105] Consequently, the structure and the interfaces of back-end microservices can remain hidden from the outside world, making the system more secure and robust. Moreover, the interactions that front-end components and external services can have with the system are restricted to the methods offered by the Gateway and implemented inside the API Gateway component as a set of interactions with the methods of back-end microservices. These endpoints are known only to the API Gateway through the ENEDI Platform's Service Registry component implemented with Consul, where each microservice registers its presence upon launch.

The ENEDI API is organized in four groups, described below:

- *Infrastructure API*, which provides access to services at the infrastructure level (data center facilities, racks, servers, virtual machines, containers, infrastructure services).
- *Monitoring API*, which gives access to the Monitoring System and its query capabilities, and allows the Dashboard to retrieve data for its visualization components.

TABLE 8 Metrics collected/calculated by ENEDI's monitoring system.

Metric name	Unit	Collection frequency	Type	Description
DC/Rack/Srv/Cont CPU utilization	%	15 s	Comp.	Total CPU utilization
DC/Rack/Srv/Cont memory utilization	%	15 s	Comp.	Total RAM utilization
DC/Rack/Srv/Cont disk I/O utilization	Kb/s	15 s	Comp.	Total Kb/s r/w on server disks
DC/Rack/Srv network utilization	Kb/s	15 s	Comp.	Total Kb/s of net. traffic to/from DC servers
Cost of kWh	€	15 s	Cost	Cost of kWh based on time period
Total energy cost of DC/Rack/Srv/Cont	€	15 s	Cost	kWh × Cost of kWh at specific period
Total energy cost savings of DC	€	15 m	Cost	Total cost saved due to solar energy utilization
A/C power consumption	W	15 s	Energy	Power of D/C cooling units
Total A/C energy consumption	kWh	15 s	Energy	Total energy consumption of D/C cooling units
Computing power consumption	W	15 s	Energy	Power of DC/Rack/Srv/Cont computing equipment
Total computing energy consumption	kWh	15 s	Energy	Total energy consumption of DC/Rack/Srv/Cont computing equipment
Total DC Power consumption	W	15 s	Energy	Total DC power consumption of IT & cooling equipment
Total DC energy consumption	kWh	15 s	Energy	Total energy consumption of of DC (computing, A/C, etc.)
Total PUE	-	15 s	Energy	Total DC energy consumption / total IT energy consumption
Total power produced by PV	kW	15 m	Energy	Power in kW produced by PV facility
Total energy produced by PV	kWh	15 m	Energy	Total energy consumption of of DC (computing, A/C, etc.)
Total CO_2 emission equivalent	g	15 m	Env.	The equivalent of grams of CO_2 emitted by the DC/Rack/Srv/Cont based on the following equation: kW/h*0.8

- *Notification API*, which gives access to system notifications and supports Websockets and 3-rd party service integration (e.g., Slack, e-mail, etc.)
- *Policy API*, which is responsible for the retrieval and manipulation of the Policies that the Decision Support System will utilize via the Policy Editor.

For the design and documentation of the API we used *Swagger*, following the *OpenAPI* Specification.[106] We did two alternative implementations of the API Gateway: one was based on *Java Spring Boot* framework[107] and the other on Javascript-based Node.js.[108]

5.4 | Platform deployment

To streamline the integration of a data center on the ENEDI platform, we developed the *ENEDI agent*, a pre-configured containerized application, which can be deployed and executed on a virtual machine of a data center to install and bootstrap the ENEDI services. The ENEDI agent runs in two modes:

- In the *data center master* mode, it is installed on a "master VM" and configures on it the containers wherein ENEDI infrastructure services are launched, namely: the Consul server for the Service Registration & Discovery Service; Prometheus as the Monitoring Service of choice; the database server for the persistent storage of configuration and log data; and the Netdata agent, which monitors the VM. These services are registered on the local Consul and on a global Consul in the case of data center federation.
- In the *data center server* mode, the agent integrates the servers of the data center's clusters by installing upon them Consul agents and monitoring probes, and registering those on the data center's Consul server.

The deployment and execution of the ENEDI agent is managed through the Configuration Widget of the Dashboard, which provides forms with parameterized configuration templates and produces Docker instructions that can install and launch the necessary services.

6 | TESTING AND EARLY EXPERIENCES

Following the implementation of the ENEDI system and the deployment of the PV facility and the platform at a University of Cyprus (UCY) data center, we conducted an analysis of collected metrics to assess the environmental and financial impact of the solar-energy adoption, and performed various benchmarks to experiment with and demonstrate the system's functionality, configurability, and performance. In this section we summarize some of these observations and tests.

6.1 | Energy savings and environmental impact

Using the data collected by the ENEDI platform from its IoT sensor infrastructure during the first year of operation, we assess the impact of improving the energy efficiency of the data center and adopting solar energy to power its operation.

Aggregate data about environmental conditions inside the data center are shown in Tables 9 and 10. Specifically, Table 9 displays the monitored environmental data (mean values) for four selected server racks during the system's first year of operation, while Table 10 depicts the monitored environmental data (mean values) for the data center rooms during the same period.

These observations comply with the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) standard regarding recommended guidelines for monitoring the environmental conditions in data centers and/or server rooms, targeting efficiency optimization.[56] In

TABLE 9 Environmental data monitored by the Indoor IoT Sensing Infrastructure for the for four selected server racks of the data center rooms during 24/09/2019 – 23/09/2020 (mean values).

Level	Point	Front temperature (°C)	Rear temperature (°C)	Front relative humidity (%)	Rear relative humidity (%)
Room 1	Rack 2	20.5	24.4	45.1	34.5
Room 1	Rack 6	20.7	25.1	50.7	33.1
Room 2	Rack 2	20.1	25.4	51.5	36.3
Room 2	Rack 6	24.3	26.3	37.5	31.3

TABLE 10 Environmental data monitored by the Indoor IoT Sensing Infrastructure for the data center rooms during 24/09/2019–23/09/2020 (mean values).

Point	Temperature (°C)	Relative humidity (%)
Room 1	21	48
Room 2	25	37

TABLE 11 Expected ENEDI PV system contribution during its first year of operation.

Month	Average daily PV production (kWh)	Expected monthly PV production (kWh)	Monetary savings (€)	CO_2 emission avoidance ($kgCO_2e$)
Jan	114,89	3,561.59	705.19	2,619
Feb	144,12	4,035.36	799.00	2,967
Mar	180,60	5,598.60	1,108.52	4,116
Apr	204,21	6,126.30	1,213.01	4,504
May	213,73	6,625.63	1,311.87	4,871
Jun	238,39	7,151.70	1,416.04	5,258
Jul	238,77	7,401.87	1,465.57	5,442
Aug	227,61	7,055.91	1,397.07	5,188
Sep	200,16	6,004.80	1,188.95	4,415
Oct	166,75	5,169.25	1,023.51	3,800
Nov	141,13	4,233.90	838.31	3,113
Dec	114,82	3.559.42	704.77	2,617
Total		**66,524.33**	**13,171.82**	**48,909**

case that any of the observed parameters exceed the ASHRAE recommendations, relevant notifications are immediately presented as alerts on the ENEDI dashboard for appropriate actions by the data center operator. The recommendations include:

- Room temperature: 18–27 °C
- Room relative humidity: <60%
- PC array temperature: 18–27 °C

The ENEDI PV system is expected to generate at least 66.52 MWh during its first year of operation, with an estimated Performance Ratio (PR) of 79.26% (considered as "very good"). As it can be seen in Table 11, the highest PV production is expected to occur during the summer season, which coincides with the period of greatest energy consumption of the data center (due to its cooling needs). Moreover, the ENEDI PV system's operation will result to the reduction of Greenhouse Gas (GHG) emissions to the environment equivalent to 48,909 $kgCO_2e$, considering a grid factor of 0.73521 $kgCO_2$ e/kWh for Cyprus as published by the Electricity Authority of Cyprus (EAC) each year and drawn by the UCY's 2019 electricity bills. During its 20-year operation, the ENEDI PV system is expected to generate more than 1200 MWh of "green" electricity, thus reducing significantly the University's carbon footprint. The PV system's contribution for the first year of operation, as simulated with the open-source PVGIS platform of the European Commission[109] is demonstrated in Table 11.

Table 12 provides a comparison between the expected monthly production of the ENEDI PV system (as obtained from the simulations) and the actual production monitored by the dedicated metering infrastructure since the system's commissioning in November 2020. As it can be seen, there is a slight difference between the two values for all months under study. In case of reduced actual values when compared to the expected ones, this difference can be considered as an effect of increased ambient temperature in relevant months, affecting PV module operation. Generally, any difference can be considered as an effect of the deviation of the PVGIS platform's data from the actual environmental conditions and PV system equipment characteristics. In any case, despite any major or minor difference observed for each distinct month, the total energy production during the 12-month period is nearly identical, that is, only about 778 kWh of difference, a value that can be considered acceptable (€158 equivalent) since it does not impact the conducted cost-benefit analysis significantly.

TABLE 12 Comparison between the expected and the actual ENEDI PV system production since its commissioning in November 2020 (for the first year of operation).

Month	Expected PV production (MWh)	Actual PV production (MWh)
Nov	4.23	3.19
Dec	3.56	3.20
Jan	3.56	3.35
Feb	4.04	4.46
Mar	5.60	5.72
Apr	6.13	6.41
May	6.63	7.43
Jun	7.15	6.94
Jul	7.40	7.16
Aug	7.06	6.71
Sep	6.00	5.95
Oct	5.17	5.22
Total	**66.5**	**65.7**

TABLE 13 Data center's energy consumption (as monitored by the dedicated metering infrastructure) breakdown and associated financial cost and environmental impact during 24/09/2019–23/09/2020.

Level	Point	Consumption (kWh)	Cost (€)	CO_2 emissions ($kgCO_2e$)
Room 1	Total server load	26,958.9	5,337.86	19, 820
	Cooling load	53,923	10,676.75	39, 645
	Total energy	80,881.9	16,014.62	59, 465
Room 2	Total server load	83,166.8	16,467.03	61, 145
	Cooling load	67,050	13,275.90	49, 296
	Total energy	150,216.8	29,742.93	110, 441
Data center	Total server load	110,125.7	21,804.89	80, 966
	Cooling load	120,973	23,952.65	88, 941
	Total energy	231,098.7	45,757.54	169, 906

Table 13 displays a breakdown of the data center's actual energy consumption (as monitored by the dedicated metering infrastructure), as well as the associated financial cost and environmental impact, during ENEDI's first year of operation at UCY. It can be seen that the server and cooling loads have nearly even shares of the total energy consumption.

As the data center's energy consumption is covered by the local utility, as well as the University's diesel generators in case of power outage, both financial and environmental savings are achieved by consuming energy produced by the PV system, instead of being purchased by the utility as a result of the operation of environmentally polluting thermal power plants.

Energy savings are both direct and indirect. Immediate energy savings result from the operation of the PV system. Reducing operating costs for energy implies the release of resources by the University for further development use. Also, the reduction of CO_2 emissions means the reduction of air pollution and the reduction of the UCY's environmental footprint. Indirect energy savings come from the operation of the information system in which the data from the metering devices are recorded, so that they can be used for interventions at the operational level as well as logistics and building infrastructure, aiming at energy savings and further reduction of the UCY's energy footprint. Moreover, the ENEDI PV system's operation increases the contribution rate regarding the UCY's production of energy from RES, its self-consumption and its self-sufficiency.

Figure 12 (left) illustrates the expected contribution of the ENEDI PV system over the total data center energy consumption during its first year of operation. It can be seen that about 30% of the total energy demand will be covered by the implemented PV system, thus resulting in a significant

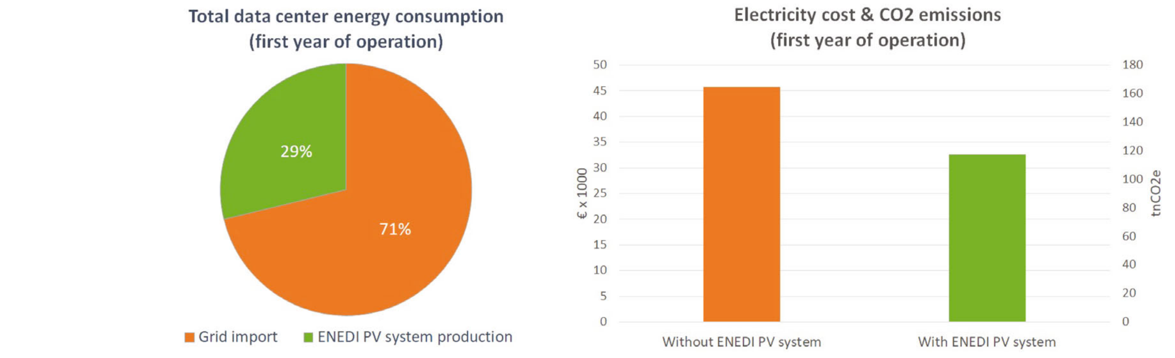

FIGURE 12 Expected contribution of the ENEDI PV system over the total data center's energy consumption during its first year of operation (left) and expected reduction in the data center's electricity cost and CO_2 emissions due to the contribution of the ENEDI PV system during its first year of operation (right).

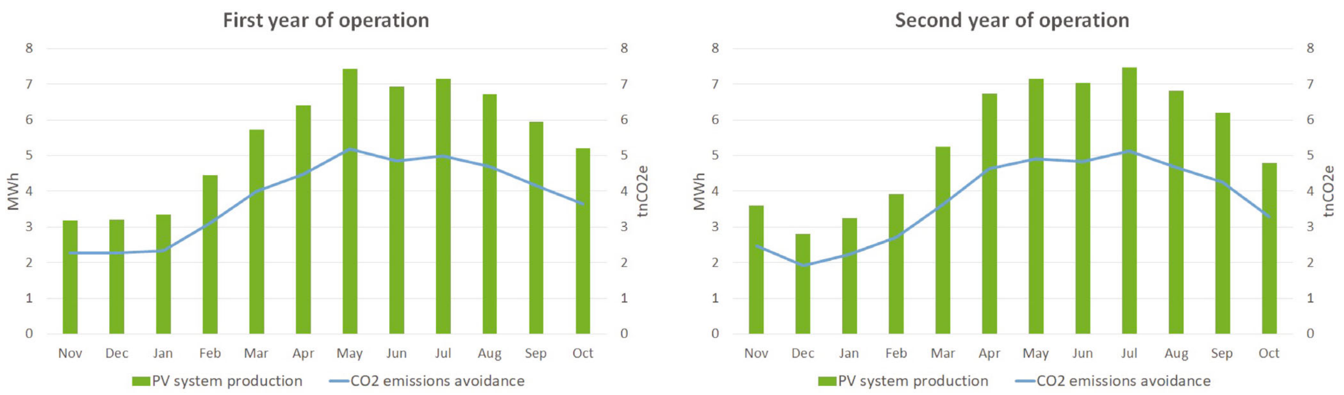

FIGURE 13 The performance of the ENEDI PV system in terms of PV production and CO_2 emissions avoidance since its commissioning in November 2020 for Year 1 (left) and Year 2 (right).

reduction in the data-center's electricity cost for that period and its environmental footprint, as indicated in Figure 12 (right). The coverage of the data center's total energy demand by 30% can be considered as a significant value and an important step towards a "green" data center, however even higher shares should be aimed. For example, a share of at least 50% could be considered as adequate in order for a data center to be regarded as "green," as at least half of its total energy consumption would be covered by clean electricity. Generally, the use of quantified metrics and predefined limits is of utmost importance, especially in the absence of a clear relevant framework for the classification of data centers with regards to their energy use requirements. To this extent, in order to reach the aforementioned 50% threshold (served as an example) of the data center's total energy use covered by clean electricity generated on-site, a 72 kWp PV system is required, instead of the implemented 41.31 kWp one. As a result, this requires the addition of approximately 31 kWp to the existing facility, thus nearly doubling the system's capacity and in turn, increasing significantly the area required for installation. Unfortunately, any increase in the ENEDI PV system's capacity is impossible, as all the available unshaded and free from obstacles (such as compressors, antennas, etc.) space on the different roofs of the data center building has already been exploited. Specifically, since the initial stages of implementation and mainly during the sizing of the PV system, the aim was to utilize all the available roof area for maximum PV system capacity and thus, maximum clean electricity generation, in an effort to minimize in turn the data canter's carbon footprint. An option would be the installation of a PV subsystem in a nearby building or any other available site within the organization, providing its output directly to the University and not to the data center itself. Moreover, for a 75% coverage, a 107 kWp PV system is necessary, while for a total energy use coverage by clean electricity, a 143 kWp PV system in needed, that is, about 3.5 times higher than the size of the implemented ENEDI PV system. It has to be noted though, that any energy efficiency and/or energy conservation measures must be taken prior to the implementation of a power generating facility, in order to avoid any unnecessary investments in power capacity, a process which was followed for the case of this work, as already explicitly explained.

The dataset obtained during the last two years of system operation, that is, since its commissioning in November 2020, demonstrates the performance of the ENEDI PV system in terms of PV production and CO_2 emissions avoidance. Figure 13 (left) illustrates the monthly PV production and the corresponding CO2 emissions avoidance for the first year of operation, while Figure 13 (right) depicts the respective values for the second

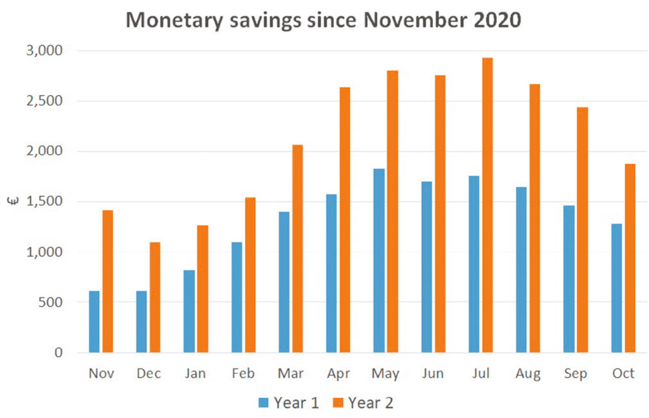

FIGURE 14 Comparison of the financial benefit of the ENEDI PV system since its commissioning in November 2020 (two years of operation).

year of operation, considering the grid factor for the relevant years (as published by the EAC) and drawn by the corresponding UCY's electricity bills. It must be noted that with the steadily increasing penetration of RES in Cyprus' energy mix per year basis, the grid factor faces a respective (slight) reduction. The results show that more than 180 MWh of clean electricity have been produced (65.75 MWh in Year 1 and 65.05 MWh in Year 2), thus resulting in saving approximately 90.7 $tnCO_2e$ (45.99 in Year 1 and 44.70 in Year 2).

In addition, the monetary savings derived by the system's operation since its commissioning are rendered in Figure 14, considering the average price per unit of electricity of the UCY as calculated according to the 2020–2022 electricity bills. It can be clearly seen that the benefit of producing on-site clean electricity from RES is enhanced even more considering the vast increase in electricity prices, which have nearly doubled when comparing (for example) the values of 2022 to the ones of 2020. Notably, the results show that the ENEDI PV system has saved more than 41k € since its commissioning in November 2020, categorized at 15,774€ in Year 1 and 25,468€ in Year 2.

Furthermore, equipping clean power generating facilities (such as the ENEDI PV system) with Energy Storage assets, specifically Battery Energy Storage Systems (BESS), can enhance even further the data center's energy sustainability. As PV systems produce clean energy during daylight, a data center's overnight energy demand is covered by the utility, essentially by conventional power generating units, which are in any case environmentally unfriendly. The integration of a BESS with the on-site PV system enables the opportunity to store any excess PV production occurred during daytime and to provide it to the data center not only during the night, but also during peak-tariff periods, thus not only reducing the cost of operation for the data center's owner, but also stabilizing the grid's operation, through the reduction of the site's grid interaction. In general, the integration of Energy Storage assets and BESS can enhance the data center's energy flexibility, with all the relevant benefits for both the owner and the grid operator. In any case though, the proper sizing of the hybrid PV-BESS system is of utmost importance as it affects significantly the financial viability of the project, especially due to the still high (but continuously falling) cost of BESS.

As the ENEDI PV system is already grid-connected, any unavailability of solar power is currently addressed by energy purchased from the grid at the retail electricity price the UCY is charged at. Similarly, any excess of generation is injected to the grid and credited to UCY for later use, due to the relevant compensation mechanism (supporting scheme for PV systems) that UCY is already benefiting from. In this way, the grid itself is utilized by the UCY as a virtual energy storage asset for the ENEDI PV system. A possible further improvement of the ENEDI PV system is, of course, the use of an actual Energy Storage asset and more specifically, in the form of a BESS (as described above), thus converting the ENEDI PV system to a hybrid PV-BESS asset, increasing in turn even further the self-consumption and self-sufficiency of the system and reducing its grid interaction. Thoughts are currently being made to this extent, which (if implemented) will also increase the system's energy flexibility and enable the availability of Demand Response actions. However, the current size of the PV system must be also considered, as PV production is presently inadequate to cover the high energy consumption of the data center to a greater extent (as the aforementioned results have showed) so that energy surplus can occur and thus, be available for storage.

6.2 | Benchmarking experiments

To demonstrate the functionality of the ENEDI platform, we ran a number of tests on a selected server of our data center, which has 16-cores and 64 GB of main memory. We kept the server idle for 60 min. During that period, ENEDI's monitoring system reports near-zero CPU utilization, 40% utilization of the RAM (attributed to Openstack services running on the background), an energy consumption of 0.5 kWh with a cost of €0.11, and the production of 140g of CO_2 (1 kWh corresponds to 0.73521 kg of CO_2, according to the DSO). Subsequently, we ran a CPU stress test, which aims at utilizing all CPU cores in full for 10min, and a RAM stress test, which aims at utilizing fully the main memory of the server by running a program

FIGURE 15 Visualization of monitoring results from CPU (left) and Main Memory stress tests (right). Screenshots are taken from a general-purpose visualization engine integrated to ENEDI.

FIGURE 16 ENEDI Dashboard Screenshots: Comparing data about PUE, power consumption, CO_2 emissions, power production, percentage of energy self-sufficiency, energy cost and energy savings of two geo-distributed data centers.

that does continuous allocations and de-allocations of 32GB of memory on a single core. The ENEDI monitoring system captures changes in the server-resource's utilization and the data-center's energy consumption, and these changes are depicted in the dashboard as they occur, along with corresponding information about the resulting energy cost and corresponding CO_2 emissions (see Figure 15).

Using Prometheus' support for hierarchical federation of Prometheus servers, we can easily arrange the collection and integration of monitoring data retrieved from geo-distributed ENEDI-compliant data centers, and its visualization and analysis on the same dashboard. To explore this capability, we provided the ENEDI agent to our partners who manage a data center at the University of Crete (UOC), Greece. This data center was configured to provide a basic set of requirements, namely servers running CoreOS Linux, Openstack, and Docker; local installation of Kubernetes; collocated PV park for solar energy production, and indoor and outdoor sensors. The installation and configuration of the ENEDI platform on UOC's data center was straightforward. Also, the federation of the two monitoring systems was easily performed via the ENEDI agent (see Section 5.4). Then, using the ENEDI dashboard's configuration capabilities, we easily configured widgets depicting data collected from both data centers and performed comparisons of relative energy costs, "green" energy use, percentage of power used that came from self-production and so forth. In Figure 16, we present screenshots from the dashboard that capture a 4-day-long benchmarking test executed simultaneously on the two data centers. In this experiment we used a synthetic load stressing 14 and 5 servers of UCY and UOC, respectively, with three 2-hour-long workloads corresponding to 40%, 60% and 80% CPU utilization, respectively.

7 | ENERGY-EFFICIENT DATA CENTERS: STATUS AND TRENDS

There is a vast research literature focusing on issues pertinent to the design, implementation, and operation of energy-efficient data centers. Recently, these issues are attracting renewed interest because of the push towards a "zero carbon" economy triggered by the perils of climate change.[12,22,30,37,57] The positive impact that renewable energy sources (RES) can bring to data center carbon emissions[16] and to the total cost of data center ownership[110] has motivated researchers to make "green" energy consumption a first-level concern in the design, placement and management of data centers.[43] This trend is further reinforced by the impressive reduction in the cost of implementing RES such as PV arrays.[58-60] Consequently, many research studies sought to analyze and understand the energy profile of data center systems[15,36,42,52,57,64,110,111] and applications,[19,25,27,33,112] their impact on the environment[10,113-115] and on the cost of owning and operating a data center.[7,61,116] Also, many industrial initiatives[35,37] are focusing on the implementation of, and experimentation with "green" data centers, which generate their own renewable energy or draw it directly from existing nearby plants. These efforts have given remarkable results, with hyper-scale data centers constructed and operated by large operators like Google, Amazon and Microsoft[18-21] expanding aggressively their share of "green" energy use[117] and attaining near-optimal Power Usage Effectiveness (PUE) values below 1.1–1.2.[1]

Numerous other research works are exploring alternative approaches to develop energy-efficient systems across all the layers of the hardware-software-application/service stack—from processors, servers, and energy management systems to applications and service pricing policies—using modeling, simulation, experiments or observation. In particular, research and industrial efforts in computer architecture are focusing on the design of energy-efficient processors, accelerators, and servers,[13] which can reduce their power consumption by switching-off completely during periods of inactivity or by switching between high and low-power states in response to fluctuating workloads,[7,64,116,118-120] provisioned power constraints ("power caps"),[121,122] or intermittent availability of "green" energy.[123]

In the area of data center middleware, many studies focus on the energy-efficient management of a data center's workload using various algorithms from machine learning,[19,124,125] self-adaptive scheduling and optimization,[126] to resource allocation, task placement and scheduling.[44,64,127] Many recent works address this problem in the context of cloud services, which are based on virtual machines or containers,[47,66,85] and modern application paradigms, such as microservices,[85,126,128] data-intensive applications,[33,52,129] and interactive web applications.[51]

Looking beyond the isolated data center, many researchers pursue "geographical load balancing," namely the synergistic allocation of tasks across cooperating geo-distributed data centers. Geographical load balancing tries to exploit spatio-temporal differences in environmental conditions, in power-grid and on-site electricity prices, and in excess computing capacity of remote data centers in order to achieve overall better reliability, increased availability, and lower end-user latency with lower cost and reduced carbon footprint.[130-133] These efforts often pursue the concept of "follow-the-sun" (for solar-powered facilities) or "follow-the-renewables" (for wind and/or solar) computation, wherein computing workloads migrate dynamically towards sites that have larger availability of renewable energy.[46,47,66,129,134-136] For example, Akoush et al. introduced a software framework that supports the migration of virtual machines between data centers according to "green" power availability,[47] whereas Qi et al. proposed scheduling algorithms for optimal VM migration.[66] Zhang et al, introduced the "GreenWare" middleware, which takes into account both the availability of renewable energy of a network of cooperating data centers and the desired cost of the Internet service operator to decide dynamically where to dispatch incoming tasks so as to maximize the total use of "green" energy.[134] In similar vein, Liu et al., proposed distributed algorithms for optimal geographical load balancing leading to a reduction in the use of energy from fossil fuel sources.[46] Yang and Chien, in,[137] proposed the "Zero-carbon Cloud" (ZCCloud) approach that is tailored to the intermittent characteristics of RES (wind, solar), which often produce "stranded" energy, namely energy going wasted during times of low computational demand. ZCCloud takes advantage of stranded "green" energy to sustain the computation of particular job types, which run on intermittent power without hurting their service-level objectives ("capability" and "on-time" jobs). Further analysis of the economic viability of ZCCloud showed that the method can bring substantial benefits to the total cost of ownership of cooperating data centers, as well as advantages for the environment and the power-grid.[31] Gupta et al. studied geographical load balancing in the context of Internet-scale Content Delivery Networks (CDN).[138] More specifically, they explored how to meet the zero-carbon emissions goal for a global network with hundreds of CDN data centers by using the minimum possible number of collocated PV panels. The authors formulate this as a Linear Programming optimization problem and propose heuristic algorithms to manage the dynamic migration of loads towards CDN data centers with higher levels of solar energy.

Other works focus on the design, implementation, and operation of "green" data centers, which are powered by collocated RES. Goiri et al.[139] presented the design and implementation of Parasol, a prototype facility comprising a small container with computing and cooling equipment, a solar panel roof and a battery bank. The authors introduced also GreenSwitch, a scheduler running on Parasol, which chooses the energy source (renewable, battery or power-grid) and the renewable energy storage medium for each particular workload. Experiments with MapReduce jobs running on Parasol and GreenSwitch demonstrated that intelligent workload and energy source management can produce significant cost reductions.[140]

A number studies seek to identify the best geographic locations for placing cooperating data centers and their collocated RES. Berral et al. formulate this as a Mixed Integer Linear Programming optimization problem[141] and propose heuristic algorithms to solve it. The authors developed also a decision-making tool for selecting data center locations according to various parameters regarding the type of RES, energy storage, desired percentage of reliance on "green" energy, and expected cost. They also introduced the GreenNebula middleware, which supports the migration of live virtual machines according to the "follow-the-renewables" principle. The placement of "green" data centers is also examined in Reference 32; in this work, however, the authors model wider demographic and macro-economic trends in developing economies of Asia that affect the long-term dynamics of data center workloads and their geographic distribution. These models can provide insights regarding decisions on when to build and where to place new data centers.

In summary, the combination of various approaches for the placement, manufacturing, RES collocation, workload management and operation of data centers can result in substantial improvements in the overall efficiency of data centers in terms of the reduction of their power consumption, carbon footprint and operational costs. However, the application of these approaches on existing medium or small data centers is not always straightforward or even possible: more often than not, such data centers are incrementally designed, procured, developed, and upgraded; their location is determined by various constraints related to institutional priorities regarding real-estate availability, budget allocation, staffing and operational support; self-production of renewable energy is regulated and subject to budget constraints and space availability.[23] Also, the federation of small and medium data centers with possibly different software and operational configurations requires a common data foundation, which would enable the sharing of usable and relevant data. In our work, we take an integral view of all these aspects and describe the regulatory, financial, and technical steps that need to be taken to transform an existing medium-size data center into a "green" one, by collocating with it a PV array, and a network of IoT sensors monitoring the operations and the energy balance of the facility. Using state-of-the-art virtualization and orchestration software, we streamline the configuration and management of the data center infrastructure, and simplify the federation of geo-distributed data centers. Using the proposed ENEDI platform, operators can easily experiment with alternative algorithms and systems for sharing workloads, migrating VMs or containers and adopt the techniques that are suitable to the needs of their end-users.

8 | CONCLUSIONS

The use of renewable energy to power small and medium data centers is attracting increased interest due to the push towards a zero-carbon economy. The transformation of conventional data centers to "green" requires, however, interventions that range from the establishment of renewable energy sources nearby the data center, to the continuous integration of data about its operational status and environmental conditions into a common data platform that makes these data shareable and amenable to actionable analytics.

In this article, we describe steps taken to transform a medium-scale academic data center into a "green" one. This transformation entailed: (i) The design, procurement and implementation of a collocated photovoltaic facility; (ii) Interventions to improve the energy-efficiency of the data center; (iii) Instrumentation of the data center's electrical, mechanical and computing equipment with IoT sensors, which monitor and disseminate performance and environmental metrics, and (iv) the design and implementation of ENEDI, a cyber-physical system for the configuration and monitoring of the infrastructure.

ENEDI manages the IoT sensors and collects metric streams regarding renewable energy production, power consumption, and environmental conditions inside and outside the data center. Also, it integrates these streams with data regarding the computing loads of its IT infrastructure and the spot price of electricity purchased from the power-grid operator. ENEDI's architecture is modular, combining the microservices paradigm with containerization, and is built on top of state-of-the-art open-source libraries. The metric streams collected by ENEDI are made available for visualization and actionable analytics through the ENEDI dashboard, which also provides support for configuring the cyber-physical system.

The platform has been deployed and is in operation. Experiments have demonstrated the functionality of the infrastructure and the dashboard, the low-cost of installing and configuring ENEDI on medium and small-scale data centers that use Openstack and Linux. Measurements from the first year of operation validate our estimates regarding the energy production and financial viability of the investment to solar energy: with an initial installation cost of nearly €53K, which was constrained by available funding and space, the photovoltaic system's net present value is estimated to €93K, its payback period to 4 years, its levelized cost of electricity to 5.5c €/kWh, which is among the lowest in Europe, and its resulting avoidance of GHG emissions to approximately 48 tnCO2e per year.

We are currently using the ENEDI platform to explore how to increase the use of renewable energy sources by small and medium-scale data centers and improve their PUE by developing: (i) decision-support systems for data-center business managers and operators regarding the management of IT and energy resources; (ii) automated workload management techniques that rely on historical and near real-time data exposed through the dashboard, and (iii) green-energy aware schemes for container migration between geo-distributed data centers. In future works, we plan to build on experiences derived from the design, implementation and deployment of, and experimentation with ENEDI, to produce valuable data sets and insights for the modeling and design of envisioned "green," energy-efficient fog infrastructures.[142] Such infrastructures will comprise micro-data centers embedded in the urban and industrial milieux, connecting through 5G or other advanced wireless technologies[143,144] to the billions of devices of the emerging IoT landscape, and supporting a variety of application paradigms, from query-driven data analytics[145] and microservices[55] to serverless.[146] Access to such data sets will enable us to explore Machine Learning-based optimization techniques for the automated configuration and run-time resource adaptation of fog and edge infrastructures and services.[147]

ACKNOWLEDGMENTS

This work was supported by the ENEDI Project, which was co-financed by the European Union (ERDF) and national funds of Greece and Cyprus (Cross-Border Cooperation Programme Greece Cyprus 2014 – 2020). The authors wish to express their gratitude to Ilias Hatzakis, Stelios Sartzetakis (GRNET – National Infrastructures for Research and Technology), Yannis Fragkiadakis, Manolis Stavrakakis, Ophelia Neofytou (University of Crete), Costas Charalambous, Savvas Nikiforou, Moysis Symeonidis, and Joanna Georgiou (University of Cyprus) for their valuable contributions that made this project possible.

CONFLICT OF INTEREST STATEMENT

All authors declare that they have no conflict of interest.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as the data produced by the system are not maintained.

ORCID

Marios D. Dikaiakos https://orcid.org/0000-0002-4350-6058

REFERENCES

1. Barroso LA, Holzle U, Ranganathan P. The Datacenter as a computer. *An Introduction to the Design of Warehouse-Scale Machines*. Vol 2. 3rd ed. Morgan & Claypool Publishers; 2018.
2. Cisco Public. *Cisco Global Cloud Index: Forecast and Methodology, 2015–2020. Technical report*. Cisco; 2016.
3. Carr N. *The Big Switch: Rewiring the World, from Edison to Google*. W. W. Norton & Company; 2013.
4. Armbrust M, Stoica I, Zaharia M, et al. A view of cloud computing. *Commun ACM*. 2010;53(4):50.
5. Ben-Yehuda OA, Ben-Yehuda M, Schuster A, Tsafrir D. The rise of RaaS. *Commun ACM*. 2014;57(7):76-84.
6. Koomey J, Brill K, Turner P. White paper: A simple model for determining true total cost of ownership for data centers. *Uptime Ins*. 2007:9.
7. Hardy D, Kleanthous M, Sideris I, Saidi AG, Ozer E, Sazeides Y. An analytical framework for estimating TCO and exploring data center design space. *ISPASS 2013 - IEEE International Symposium on Performance Analysis of Systems and Software*. IEEE Computer Society; 2013:54-63.
8. Masanet E, Shehabi A, Lei N, Smith S, Koomey J. Recalibrating global data center energy-use estimates: growth in energy use has slowed owing to efficiency gains that smart policies can help maintain in the near term. *Science*. 2020;367(6481):984-986.
9. Koomey J. *Growth in Data Center Electricity Use 2005 to 2010*. Analytics Press; 2011.
10. Jones N. How to stop data centres from gobbling up the world's electricity. *Nature*. 2018;561:163-166.
11. Andrae A, Edler T. On global electricity usage of communication technology: trends to 2030. *Challenges*. 2015;6(1):117-157.
12. Koronen C, Åhman M, Nilsson LJ. Data centres in future European energy systems–energy efficiency, integration and policy. *Energ Effic*. 2020;13(1):129-144.
13. Hennessy JL, Patterson DA. A new golden age for computer architecture. *Commun ACM*. 2019;62(2):48-60.
14. Woods A. Cooling the data center. *Queue*. 2010;8(3):1-10.
15. Mastelic T, Oleksiak A, Claussen H, Brandic I, Pierson JM, Vasilakos AV. Cloud computing: survey on energy efficiency. *ACM Comput Surv*. 2015;47(2):1-36.
16. Frazelle J. Power to the people. Reducing datacenter carbon footprints. *ACM Queue*. 2020;18(2):47.
17. Shehabi A, Smith S, Sartor D, et al. *United States Data Center Energy Usage Report*. Technical Report LBNL-1005775, Lawrence Berkeley National Lab; 2016.
18. Google. *Efficiency – Data Centers*, 2021.
19. Gao J, Jamidar R. Machine learning applications for data center optimization. *Google White Paper*. 2014;1-13.
20. Gamble C, Gao J. *Safety-First AI for Autonomous Data Centre Cooling and Industrial Control*; 2018.
21. Smith B, Browne CA. *Tools and Weapons: the Promise and the Peril of the Digital Age*. Penguin Press; 2019.
22. Chien AA. Driving the cloud to true zero carbon. *Commun ACM*. 2021;64(2):5.
23. Vasques TL, Moura P, de Almeida A. A review on energy efficiency and demand response with focus on small and medium data centers. *Energ Effic*. 2019;12(5):1399-1428.
24. Feldmann A, Gasser O, Lichtblau F, et al. A year in lockdown. *Commun ACM*. 2021;64(7):101-108.
25. Patterson D, Gonzalez J, Le Q, et al. *Carbon Emissions and Large Neural Network Training. arXiv*, apr. 2021.
26. Mathis W, Saul J, Cang A, Huang Z. Bitcoin miners navigate extreme world of crypto power-hunting. *Bloomberg*. Cambridge Centre for Alternative Energy; 2021.
27. Cambridge Centre for Alternative Finance. *Cambridge Bitcoin Electricity Consumption Index*, 2021.
28. Kamiya G, Kvarnstrom O. *Data Centres and Energy – from Global Headlines to Local Headaches? Technical report*. International Energy Agency (IEA); 2019.
29. Montevecchi F, Stickler T, Hintemann R, Hinterholzer S. *Energy-Efficient Cloud Computing Technologies and Policies for an Eco-Friendly Cloud Market*. European Union; 2020.
30. Larson E, Greig C, Jenkins J, et al. *Net-Zero America: Potential Pathways, Infrastructure, and Impacts Interim Report. Technical report*. Princeton University; 2020.
31. Yang F, Chien AA. Large-scale and extreme-scale computing with stranded green power: Opportunities and costs. *IEEE Trans Parallel Distribut Syst*. 2018;29(5):1103-1116. doi:10.1109/TPDS.2017.2782677
32. Liu R, Sun W, Weisheng H. Workload based geo-distributed data Center planning in fast developing economies. *IEEE Access*. 2020;8:224269-224282.
33. Goiri Í, Le K, Nguyen TD, Guitart J, Torres J, Bianchini R. GreenHadoop: leveraging green energy in data-processing frameworks. *EuroSys'12 - Proceedings of the EuroSys 2012 Conference*. ACM; Vol 57–70:2012.
34. The Center of Expertize for Energy Efficiency in Data Centers U.S. Dept. of Energy. Federal Energy Management Program, 2020.
35. Google. *Renewable Energy*, 2020.
36. Amazon. *Sustainability in the Cloud*, 2020.
37. Microsoft. *Transitioning to zero-carbon energy*, 2020.
38. Gloukhovtsev M. *How 5G Transforms Cloud Computing. Technical report*. Dell Technologies; 2020.
39. El-Sayed H, Sankar S, Prasad M, et al. Edge of things: the big picture on the integration of edge, IoT and the cloud in a distributed computing environment. *IEEE Access*. 2017;6:1706-1717.
40. Martinez I, Hafid AS, Jarray A. Design, resource management and evaluation of fog computing systems: A survey. *IEEE Internet Things J*. 2021;8(4):2494-2516. doi:10.1109/JIOT.2020.3022699
41. White MC. Our digital lives drive a brick-and-mortar boom in data centers. *New York Times*. 2020.
42. Avgerinou M, Bertoldi P, Castellazzi L. Trends in data Centre energy consumption under the European code of conduct for data Centre energy efficiency. *Energies*. 2017;10:18.

43. Stewart C, Shen K. Some joules are more precious than others: Managing renewable energy in the datacenter. *Proceedings of the Workshop on Power Aware Computing and Systems (HotPower'09)*. USENIX Association; 2009.
44. Mazzucco M, Dyachuk D, Dikaiakos MD. Profit-aware server allocation for green internet services. *Proceedings - 18th Annual IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems. MASCOTS*; 2010:2010.
45. Liu Z, Lin M, Wierman A, Low SH, Andrew LLH. Geographical load balancing with renewables. *Performance Evaluat Rev*. 2011;39(3):62–66.
46. Liu Z, Lin M, Wierman A, Low S, Andrew LLH. Greening geographical load balancing. *IEEE/ACM Trans Network*. USENIX Association; 2015.
47. Akoush S, Sohan R, Rice A, Moore AW, Hopper A. Free lunch: Exploiting renewable energy for computing. *13th Workshop on Hot Topics in Operating Systems, HotOS 2011*. ACM; 2011.
48. Kong F, Liu X. A survey on green-energy-aware power management for datacenters. *ACM Comput Surv*. 2014;47(2):1-38.
49. Hasan MS. *Smart Management of Renewable Energy in Clouds: From Infrastructure to Application*. PhD thesis, INSA de Rennes; 2017.
50. Hasan MS, Kouki Y, Ledoux T, Pazat JL. Exploiting renewable sources: When green SLA becomes a possible reality in cloud computing. *IEEE Trans Cloud Comput Secur*. 2017;5(2):249-262.
51. Toosi AN, Chenhao Q, de Assunção DM, Buyya R. Renewable-aware geographical load balancing of web applications for sustainable data centers. *J Netw Comput Appl*. 2017;83:155-168.
52. Wang L, Khan SU. Review of performance metrics for green data centers: A taxonomy study. *J Supercomput*. 2013;63(3):639-656.
53. Haque ME, Le K, Goiri I, Bianchini R, Nguyen TD. Providing green SLAs in high performance computing clouds. *2013 International Green Computing Conference Proceedings, IGCC 2013*. IEEE Computer Society; 2013.
54. Fowler SJ. *Production-Ready Microservices*. IEEE Computer Society; 2017.
55. Trihinas D, Tryfonos A, Dikaiakos M, Pallis G. DevOps as a service: Pushing the boundaries of microservice adoption. *IEEE Internet Comput*. 2018;22(3):65-71.
56. ASHRAE *Technical Committee 9.9. Data Center Power Equipment Thermal Guidelines and Best Practices. Technical report*. American Society of Heating, Refrigerating and Air-Conditioning Engineers; 2016.
57. Shuja J, Gani A, Shamshirband S, Ahmad RW, Bilal K. Sustainable cloud data Centers: a survey of enabling techniques and technologies. *Renew Sust Energ Rev*. 2016;62:195-214.
58. International Energy Agency. *Evolution of Solar PV Module Cost by Data Source*. International Energy Agency; 2021:1970-2020.
59. Barzegkar-Ntovom GA, Chatzigeorgiou NG, Nousdilis AI, et al. Assessing the viability of battery energy storage systems coupled with photovoltaics under a pure self-consumption scheme. *Renew Energy*. 2020;152(C):1302-1309.
60. Jäger-Waldau A. *PV Status Report 2019. Technical Report*. European Commission, Joint Research Centre; 2019.
61. Gürhan Kök A, Shang K, Yücelc S. Impact of electricity pricing policies on renewable energy investments and carbon emissions. *Manag Sci*. 2018;64(1):131-148.
62. Wikipedia. Financial incentives for photovoltaics, 2021.
63. *Electricity Authority of Cyprus*. Net Billing; 2021.
64. Hsu CH, Deng Q, Mars J, Tang L. SmoothOperator: Reducing power fragmentation and improving power utilization in large-scale datacenters. *ACM SIGPLAN Notices*. Vol 53. Association for Computing Machinery; 2018:535-548.
65. Hsu YF, Kuwahara H, Matsuda K, Matsuoka M. Toward a workload allocation optimizer for power saving in data centers. *Proceedings - 2019 IEEE International Conference on Cloud Engineering, IC2E 2019*. Institute of Electrical and Electronics Engineers Inc.; 2019:56-66.
66. Qi L, Chen Y, Yuan Y, Shucun F, Zhang X, Xiaolong X. A QoS-aware virtual machine scheduling method for energy conservation in cloud-based cyber-physical systems. *World Wide Web*. 2020;23(2):1275-1297.
67. Chatzigeorgiou NG, Poize N, Florides MA, Georghiou GE. Analysing the operation of residential photovoltaic – battery storage Systems in Cyprus. *The 12th Mediterranean Conference on Power Generation, Transmission, Distribution and Energy Conversion (MEDPOWER2020)*; 2020.
68. VDE. *Power Generating Plants in the Low Voltage Grid*Technical report, VDE VERLAG GMBH; 2019.
69. Canonical. Metal-as-a-service. https://maas.io/, 2022.
70. RedHat. Kernel virtual machine. https://www.linux-kvm.org/page/Main_Page.
71. Vogels W. Beyond server consolidation. *Queue*. 2008;6(1):20-26.
72. Pietri I, Sakellariou R. Mapping virtual machines onto physical Machines in Cloud Computing. *ACM Comput Surv*. 2016;49(3):1-30.
73. Rackspace Technology. Openstack. https://www.openstack.org.
74. Lamourine M. OpenStack. *Log*. 2014;39(3).
75. Docker. https://www.docker.com, 2022.
76. Bernstein D. Containers and cloud: from LXC to docker to kubernetes. *IEEE Cloud Comput*. 2014;1(3):81-84.
77. Sharma P, Chaufournier L, Shenoy P, Tay YC. Containers and virtual machines at scale: a comparative study. *Proceedings of the 17th International Middleware Conference, Middleware*. Association for Computing Machinery, Inc.; 2016:1-13.
78. Randal A. The ideal versus the real: revisiting the history of virtual machines and containers. *ACM Comput Surv*. 2020;53(1):1-31.
79. Laurie D. Ipmitool. https://github.com/ipmitool/ipmitool.
80. Netdata. Netdata monitoring tool.
81. Intel Hewlett-Packard NEC Dell. Intelligent Platform Management Interface Specification, v2.0 rev. 1.1, October.
82. Li C, Xue Y, Wang J, Zhang W, Li T. Edge-oriented computing paradigms: A survey on architecture design and system management. *ACM Comput Surv*. 2018;51(2):1-34.
83. Ngu AH, Gutierrez M, Metsis V, Nepal S, Sheng QZ. Iot middleware: A survey on issues and enabling technologies. *IEEE Internet Things J*. 2016;4(1):1-20.
84. Da Cruz MAA, Joel José PC, Rodrigues JA-M, Korotaev VV, Victor HC, Albuquerque D. A reference model for internet of things middleware. *IEEE Internet Things J*. 2018;5(2):871-883.
85. Hou X, Liu J, Li C, Guo M. Unleashing the scalability potential of power-constrained data center in the microservice era. *48th International Conference on Parallel Processing (ICPP2019)*. ACM; 2019:10.
86. Kang H, Le M, Tao S. Container and microservice driven design for cloud infrastructure DevOps. *Proceedings - 2016 IEEE International Conference on Cloud Engineering, IC2E 2016*. IEEE Computer Society; 2016:202-211.
87. Trihinas D, Tryfonos A, Dikaiakos M. Designing Scalable and Secure Microservices by Embracing DevOps-as-a-Service Offerings. 2018.

88. Zimmermann O. Microservices tenets: Agile approach to service development and deployment. *Comput Sci Res Dev*. 2017;32(3-4):301-310.
89. Balalaie A, Heydarnoori A, Jamshidi P. Microservices architecture enables DevOps: Migration to a cloud-native architecture. *IEEE Softw*. 2016;33(3):42-52.
90. Leite L, Rocha C, Kon F, Milojicic D, Meirelles P. A survey of DevOps concepts and challenges. *ACM Comput Surv*. 2019;52(6):1-35.
91. Wikipedia. Single Page Application, 2021.
92. Angular Web Framework. https://angular.io, 2021.
93. NGINX. https://www.nginx.com, 2021.
94. Kubernetes, 2021.
95. Rancher, 2021.
96. HashiCorp. Consul . https://www.consul.io/, 2021.
97. WebSocket. Intelligent Platform Management Interface.
98. socket.io 2021.
99. Apache kafka, 2021.
100. Trihinas D, Pallis G, Dikaiakos MD, Catascopia J. Monitoring elastically adaptive applications in the cloud. *Proceedings - 14th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing, CCGrid 2014* . IEEE Computer Society; 2014.
101. Trihinas D, Pallis G, Dikaiakos MD. Admin: Adaptive monitoring dissemination for the internet of things. *2017 IEEE Conference on Computer Communications INFOCOM 2017*. IEEE; 2017:1-9.
102. Trihinas D, Pallis G, Dikaiakos MD. Low-cost adaptive monitoring techniques for the internet of things. *IEEE Trans Serv Comput*. 2021;14(2):487-501.
103. Prometheus. 2019.
104. Google cloud's operations suite (formerly stackdriver), 2021.
105. Pautasso C, Zimmermann O, Leymann F. Restful web services vs. big web services: Making the right architectural decision. *In 17th International World Wide Web Conference (WWW2008)*. ACM; 2008:805-814.
106. *API Documentation and Design Tools for Teams*. Swagger; 2021.
107. Spring Boot, 2021.
108. node.js, 2021.
109. PVGIS. Photovoltaic Geographical Information System, 2019.
110. Masanet ER, Brown RE, Shehabi A, Koomey JG, Nordman B. Estimating the energy use and efficiency potential of U.S. data centers. *Proceedings of the IEEE*. Vol 99. Institute of Electrical and Electronics Engineers Inc.; 2011:1440-1453.
111. Dayarathna M, Wen Y, Fan R. Data center energy consumption modeling: A survey. *IEEE Commun Surv Tutorials*. 2016:732-794.
112. Lacoste A, Luccioni A, Schmidt V, Dandres T. Quantifying the carbon emissions of machine learning. 2019.
113. Podder S, Burden A, Singh KS, Maruca R. How green is your software? *Harv Bus Rev*. 2020:9.
114. Thopson C. How 'sustainable' web design can help fight climate change. *Wired*. 2020.
115. Taina J. How green is your software? In Pasi Tyrväinen. In: Jansen S, Cusumano MA, eds. *Software Business*. Springer Berlin Heidelberg; 2010:151-162.
116. Nikolaou P, Sazeides Y, Lampropoulos A, et al. On the evaluation of the Total-cost-of-ownership trade-offs in edge vs cloud deployments: A wireless-denial-of-service case study. *IEEE Trans Sustain Comput*. 2022;7(2):334-345 .
117. Oberhaus D. Amazon, Google, Microsoft: Here's who has the greenest cloud WIRED. *Wired*. 2019.
118. Meisner D, Gold BT, Wenisch TF. PowerNap: Eliminating server idle power. *International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS*. ACM; 2009.
119. AlLee G. Green microprocessor and server design. *Data Center Handbook*. John Wiley & Sons, Inc; 2014:401-418.
120. Hou X, Hao L, Li C, Chen Q, Zheng W, Guo M. Power grab in aggressively provisioned data Centers: What is the risk and what can Be done about it. *In 2018 IEEE 36th International Conference on Computer Design (ICCD)*. IEEE; 2018:26-34.
121. Shen K, Shriraman A, Dwarkadas S, Zhang X, Chen Z. Power containers: An OS facility for fine-grained power and energy management on multicore servers. *ACM SIGPLAN Notices*. Vol 48. ACM Press; 2013:65-76.
122. Azimi R, Badiei M, Zhan X, Li N, Reda S. Fast decentralized power capping for server clusters. *Proceedings - International Symposium on High-Performance Computer Architecture*. IEEE Computer Society; 2017.
123. Sharma N, Krishnappa D, Barker S, Irwin D, Shenoy P. Managing server clusters on intermittent power. *PeerJ Comput Sci*. 2015;1:e34.
124. Zhong Z, Minxian X, Rodriguez MA, Chengzhong X, Buyya R. Machine learning-based orchestration of containers: A taxonomy and future directions. *ACM Comput Surv*. 2022;54(10s):1-35.
125. Minxian X, Toosi AN, Bahrani B, Razzaghi R, Singh M. Optimized renewable energy use in green cloud data Centers. In: Yangui S, Rodriguez IB, Drira K, Tari Z, eds. *Service-Oriented Computing*. Lecture Notes in Computer Science. Vol 11895. Springer International Publishing; 2019:314-330.
126. Minxian X, Toosi AN, Buyya R. A self-adaptive approach for managing applications and harnessing renewable energy for sustainable cloud computing. *IEEE Trans Sustain Comput*. 2021;6(4):544-558.
127. Beloglazov A, Abawajy J, Buyya R. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Futur Gener Comput Syst*. 2012;28(5):755-768.
128. Minxian X, Buyya R. BrownoutCon: A software system based on brownout and containers for energy-efficient cloud computing. *J Syst Softw*. 2019;155:91-103.
129. Lin G, Zeng D, Li P, Guo S. Cost minimization for big data processing in geo-distributed data centers. *IEEE Trans Emerg Top Comput*. 2014;2(3):314-323.
130. Greenberg A, Hamilton J, Maltz DA, Patel P. The cost of a cloud. *ACM SIGCOMM Comput Commun Rev*. 2009;39(1):68-73.
131. Abdo JB, Demerjian J, Chaouchi H, Barbar K, Pujolle G. Cloud federation means cash. *The Third International Conference on e-Technologies and Networks for Development (ICeND2014)*. IEEE; 2014:39-42.
132. Rahman A, Liu X, Kong F. A survey on geographic load balancing based data center power management in the smart grid environment. *IEEE Commun Surv Tutorials*. 2014;16(1):214-233.
133. Smpokos G, Elshatshat MA, Lioumpas A, Iliopoulos I. On the energy consumption forecasting of data Centers based on weather conditions: Remote sensing and machine learning approach. *11th International Symposium on Communication Systems, Networks and Digital Signal Processing, CSNDSP 2018*. Institute of Electrical and Electronics Engineers Inc.; 2018.

134. Zhang Y, Wang Y, Wang X. GreenWare: Greening cloud-scale data centers to maximize the use of renewable energy. *ACM/IFIP/USENIX International Conference on Distributed Systems Platforms and Open Distributed Processing*. Springer; 2011:143-164.
135. Khosravi A, Garg SK, Buyya R. Energy and carbon-efficient placement of virtual machines in distributed cloud data centers. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Springer; 2013.
136. Hogade NS, Pasricha S, Siegel HJ. Energy and network aware workload Management for Geographically Distributed Data Centers. *IEEE Transactions on Sustainable Computing*. IEEE; 2021:1.
137. Yang F, Chien AA. ZCCloud: Exploring wasted green power for high-performance computing. *IEEE 30th International Parallel and Distributed Processing Symposium*. IPDPS; 2016.
138. Gupta V, Shenoy P, Sitaraman RK. Efficient solar provisioning for net-zero internet-scale distributed networks. *In 2018 10th International Conference on Communication Systems and Networks, COMSNETS 2018*. Institute of Electrical and Electronics Engineers Inc.; 2018:372-379.
139. Goiri Í, Katsak W, Le K, Nguyen TD, Bianchini R. Parasol and GreenSwitch. *Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '13*. ACM Press; 2013:51.
140. Goiri I, Katsak W, Le K, Nguyen TD, Bianchini R. Designing and managing data centers powered by renewable energy. *IEEE Micro*. 2014;34(3):8-16.
141. Berral JL, Goiri I, Nguyen TD, Gavalda R, Torres J, Bianchini R. Building green cloud Services at Low Cost. *In 2014 IEEE 34th International Conference on Distributed Computing Systems*. IEEE; 2014:449-460.
142. Symeonides M, Georgiou Z, Trihinas D, Pallis G, Dikaiakos MD. Fogify: A fog computing emulation framework. *5th IEEE/ACM Symposium on Edge Computing, SEC*. IEEE; 2020:42-54.
143. Symeonides M, Trihinas D, Pallis G, Dikaiakos MD, Psomas C, Krikidis I. 5g-slicer: An emulator for mobile IoT applications deployed over 5g network slices. *Seventh IEEE/ACM International Conference on Internet-of-Things Design and Implementation, IoTDI*. IEEE; 2022:115-127.
144. Symeonides M, Trihinas D, Pallis G, Dikaiakos MD. Demo: Emulating 5g-ready mobile iot services. *Seventh IEEE/ACM International Conference on Internet-of-Things Design and Implementation, IoTDI*. Vol 2022. IEEE; 2022:113-114.
145. Georgiou Z, Symeonides M, Trihinas D, Pallis G, Dikaiakos MD. Streamsight: A query-driven framework for streaming analytics in edge computing. In: Sill A, Spillner J, eds. *11th IEEE/ACM International Conference on Utility and Cloud Computing, UCC*. Vol 2018. IEEE Computer Society; 2018:143-152.
146. Castro P, Ishakian V, Muthusamy V, Slominski A. The rise of serverless computing. *Commun ACM*. 2019;62(12):44-54.
147. Gill SS, Minxian X, Ottaviani C, et al. Ai for next generation computing: Emerging trends and future directions. *Internet Things*. 2022;19:100514.

How to cite this article: Dikaiakos MD, Chatzigeorgiou NG, Tryfonos A, et al. A cyber-physical management system for medium-scale solar-powered data centers. *Concurrency Computat Pract Exper*. 2023;35(10):e7658. doi: 10.1002/cpe.7658