ENEDI: Energy Saving in Datacenters

Athanasios Tryfonos, Andreas Andreou, Nicholas Loulloudes
George Pallis, Marios D. Dikaiakos
Dept. of Computer Science
University of Cyprus
Nicosia, Cyprus
{a.tryfonos, aandre28, loulloudes.n, gpallis, mdd}@cs.ucy.ac.cy

Nikolas Chatzigeorgiou, George E. Georghiou
Dept. of Electrical and Computer Engineering
University of Cyprus
Nicosia, Cyprus
{nchatz05, geg}@ucy.ac.cy

***Abstract*—Despite significant advances in the design and fabrication of power-efficient electronics and microprocessors, and the ubiquitous availability of Class A-rated lightning, cooling and heating appliances, modern datacenters account for a growing percentage of global energy consumption, currently estimated at 1.5%. The effort to reduce the environmental impact and carbon footprint of datacenter operations has led to the emergence of *green datacenters* designed to reduce energy consumption and/or use renewable energy to power computing and peripheral devices. One of the challenges that arise in the context of solar-powered datacenters is the intelligent adjustment of energy consumption to variations of solar energy production, changing environmental conditions (temperature, humidity, etc.), and fluctuations in computing demand. In this paper, we present the design and early implementation of ENEDI, an integrated system that collects, integrates and analyzes data from Internet-of-Things (IoT) sensors, operating system, as well as, cloud middleware monitors, and open-data sources regarding: i) photovoltaic PV energy production in solar-powered datacenter facilities; ii) weather conditions; iii) power consumption by datacenter subsystems (clusters, cluster nodes, and cooling equipment), and iv) the time-evolving profiles of containerized cloud applications running on the datacenter.**

***Index Terms*—datacenter, photovoltaics, monitoring system, sensors, workload management, green computing, cloud, IoT**

I. Introduction

Cloud computing adoption has soared over the past decade due to exciting business opportunities and revenue potential. The ever-rising demand for virtualized resources has led to the establishment of large-scale datacenters that host hundreds to many thousands of physical servers. Such infrastructures draw significant amounts of electricity, leading to extensive operational costs. Despite remarkable advances in the design and fabrication of power-efficient electronics and microprocessors, as well as the ubiquitous availability of Class A-rated cooling and heating equipment, modern datacenters account for approximately 1.5% of the global energy consumption [1]. Moreover, the production of electricity used to power such datacenters relies predominately on carbon-intensive fossil fuels. Consequently, datacenter operations exhibit an increasing contribution to the worldwide carbon footprint and bring a substantial toll on our environment [1]. To minimize operating costs and mitigate carbon emissions, pioneering cloud providers including Google [2], Amazon [3] and Microsoft [4], are leading initiatives to promote the concept of ”Green Datacenters”, designed and developed to be powered 100% with clean and renewable energy [5]. Switching the energy-production that is used to power existing datacenters from fossil-fuel to solar entails a number of technical and regulatory hurdles, and raises a number of interesting trade-offs related to the potential operational, environmental and financial impact of using solar versus traditional energy sources.

In this paper, we present early experiences from the ENEDI project [6], which aims at improving the energy-efficiency of datacenters through a combination of: i) PV system deployment on datacenter buildings; ii) the improvement of datacenter building energy efficiency, and iii) the use of an intelligent workload management system (WMS), which performs computational resource provisioning by considering variations in solar energy production across datacenters.

In the following sections, we present the high-level architecture and early implementation of the ENEDI platform. ENEDI is a multi-layered platform that adheres to the architectural principles of the microservices paradigm. The main layers of the platform consist of: i) an IoT sensor layer for measuring performance, efficiency, energy production and environmental metrics of PV panels; ii) an IoT sensor layer for measuring operational parameters of computational and peripheral datacenter equipment; iii) a federated monitoring system for collecting the aforementioned infrastructure measurements and mashing them with application-level performance indicators; and iv) a decision support system that enables real-time deployments of *green applications* based on user-defined policies and metrics.

ENEDI's monitoring system collects, integrates and analyzes data from IoT sensors, operating system, and cloud middleware monitors, and open-data sources about: i) PV energy production in solar-powered datacenter facilities; ii) weather conditions; iii) power consumption by datacenter subsystems (clusters, cluster nodes, and cooling equipment), and iv) the time-evolving profiles of containerized cloud applications running on the datacenter mash them with application level indicators and feed them to its decision support system that enables real-time deployments of green applications based on user-defined policies.

The focal point of our work is to address the challenges arising from adopting solar powered datacenters by introducing ENEDI. Specifically, the contributions of this work are:

- A general public web-based responsive dashboard, that aims to raise awareness regarding the impact of using

renewable energy sources to power datacenters. Through intuitive statistical visualizations of operational parameters collected from the ENEDI monitoring system, the dashboard will highlight the positive effects of the ENEDI ecosystem in terms of cost and power efficiency.
- A collection of open-source monitoring probes, for the real-time collection of metrics originating from the PV data logger devices that will be provided to the community. Data loggers are vendor-specific specialized devices that store and process analog data using proprietary software. ENEDI allows to extract data from such closed-source devices through open-source probing interfaces. Not only will those probes operate under the ENEDI Monitoring System, but collected data could be provided as inputs to other widely adopted monitoring systems.
- A proposition of best practices that will guide future Geo-Distributed, Multi-Datacenter, Workload Management Systems to facilitate dynamic resource provisioning and allocation across federated green datacenters offering public cloud services.

The rest of the paper is structured as follows: Section II presents related work in the fields of Monitoring Systems and Green Datacenters. Section III presents in detail the high-level architecture and design decisions of the ENEDI WMS. Finally, Section IV concludes the paper, presents the current efforts in the development of ENEDI WMS and outlines the future work.

II. Related Work

This section presents the latest technological progress in the fields of Monitoring Systems and green datacenters. ENEDI will utilize advancements from these domains to realize its smart WMS. In specific, we consider monitoring solutions both at the application and infrastructure level. Furthermore, we investigate research work in the area of *green datacenters*.

A. Monitoring Systems

Monitoring is of fundamental significance in the ENEDI ecosystem. It applies to a broad spectrum, varying from cloud application monitoring to infrastructure performance measurements. Proprietary monitoring solutions such as Amazon CloudWatch [7], CloudMonix [8], and RackSpace CloudKick [9] provide Monitoring-as-a-Service to Cloud service consumers. Despite their intuitive operation, full documentation and well-integration with the underlying platform, their significant drawback is their commercial and proprietary nature which limits them to operating on specific Cloud providers.

General-purpose monitoring tools such as Ganglia [10] and Nagios [11] are used traditionally by system administrators to monitor slowly changing or static large-scale distributed infrastructures. This limitation makes these tools inadequate for addressing a swiftly adapting and dynamic cloud infrastructure. Prometheus [12], another popular monitoring system, uses mainly polling techniques to collect metrics from services. The strong points for the pull approach are that there is no need to install an agent and that the metrics can be pulled by multiple Prometheus instances. On the weak side, Prometheus does not offer durable long-term storage and scaling becomes an issue in large deployments. Finally, JCatascopia [13] is a multi-layer, interoperable Cloud Monitoring System that follows an agent-based producer-consumer approach, which provides scalable, real-time collection of metrics from multiple layers of the underlying infrastructure and deployed applications. During the metrics collection process, JCatascopia takes into consideration the rapid changes that occur due to the enforcement of elasticity actions on the application execution environment and the Cloud infrastructure.

B. Green Datacenters

Green datacenters have been thoroughly studied in recent years. Efforts focus on reducing the power bill of large-scale datacenters by combining renewable energy sources with workload scheduling and placement techniques. Authors in [14] propose a parallel job scheduler for data-centers powered partially by solar panels. Their proposed system, GreenSlot, seeks to maximize the green energy consumption while meeting the jobs' deadlines. GreenSlot predicts the amount of solar energy that will likely be available in the future, using historical data and weather forecasts. Based on predictions and the information provided by users, it schedules the workload by creating resource reservations for the future. Similarly, Goiri et al. [15] propose GreenHadoop, a MapReduce [17] framework for PV-powered datacenters. GreenHadoop predicts the amount of solar energy that will be available in the near future and schedules MapReduce jobs to maximize the green energy consumption within job time bounds. The main limitation of both works is that they operate on single location datacenters. ENEDI's WMS will operate on Geo-Distributed datacenters and will be able to schedule workload execution on time slots and location that maximize green energy utilization.

Work is also done in adapting routing policies to optimize the cost, emissions, and/or renewable energy utilization for geographically distributed datacenters. Liu et. al. [4] describe an algorithm for geographical load balancing with the goal of cost minimization. They discuss the possibility of powering an Internet-scale datacenter system entirely using on-site renewable energy sources and show that geographical load balancing increases the required capacity of renewable energy.

C. Beyond the State-of-the-Art

In contrast to previous works, ENEDI focuses on monitoring real-time data concerning actual energy production from solar panels installed on datacenter buildings, dynamic energy consumption required to run and cool datacenter equipment (clusters, cluster nodes, A/C units), dynamic resource utilization of computing resources, and open-data regarding weather conditions and weather predictions (sunshine, clouds, outside temperature, etc.). These data are collected thanks to a cyber-physical infrastructure installed inside the datacenter and which feeds an integrated monitoring system that facilitates the federation of geo-distributed datacenters. In particular, federated ENEDI datacenters adhere to the "Follow The Sun Computation" model. Based on this model, the computational

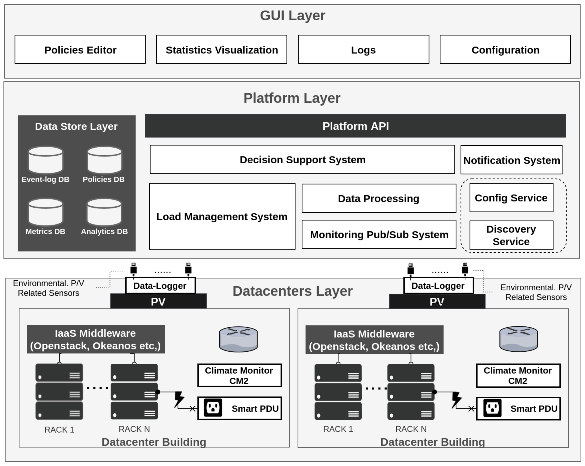

Fig. 1: ENEDI High-Level Architecture

workload in the form of stateless Docker Containers executes at time slots and servers located in datacenters that maximize green energy utilization due to energy produced by solar panels and reduce operational costs.

III. System Architecture

This section introduces the reference architecture of the ENEDI platform. It is established on the principles of the microservices, a design paradigm that structs applications as a collection of loosely coupled services that implement distinct business capabilities [18]. Figure 1 illustrates the ENEDI architecture and its components that can be grouped into three layers. Below we provide an overview of each layer.

Infrastructure Layer is the physical compartment that hosts all necessary computational equipment and support peripherals. Computational equipment consists of bare-metal servers (processing units), storage, and various inter-connectivity devices that are core in facilitating cloud operations. Support peripherals refer to integrated systems such as the power and cooling infrastructure, as well as, various sensors that continuously observe and measure an extended collection of performance, environmental and power production/consumption parameters. In the context of the ENEDI model, datacenter energy needs are satisfied via a hybrid model of renewable and conventional Grid power. Renewable energy is provided via modern photovoltaic(PV) arrays deployed on the rooftop of the datacenter building or in nearby facilities.

Platform Layer is the responsible entity to facilitate the concept of *Green computing*, established via two fundamental processes: (a) *Monitoring* and (b) *Analysis and Decision Making*. The former collects fine-grained measurements from both, the underlying infrastructure and running applications following a user-defined sampling frequency. Subsequently all metrics are stored and tagged into a time-series database, readily available for request and dissemination to other system components. The Analysis and Decision Making process digests the collected metrics, applies data analysis techniques. By taking into consideration user-defined energy policies and by implementing state-of-the-art algorithms (GreenSlot, Greenhadoop and Green Load Balancing), will support datacenter administrators in the decision making process whether and how the transfer of computational workload (load-balancing) will take place between interconnected datacenters.

GUI Layer; The graphical interface will enable end-users to interact with the underlying layers of the system. Specifically, users will perform fine-grained configurations and tuning of the various ENEDI components. Such configurations include the adjustment of monitoring frequency, the selection of the notifications medium and the definition of logging events among others. Through the *Policies Editor* component, administrators will have the ability to define datacenter energy related policies, based on which operational costs and resource utilization will be improved. Most importantly, the GUI will provide intuitive visualizations and plots that will showcase the impact of the ENEDI platform on the green transformation of datacenters.

In the following subsections, we dive deeper into the specifics of the Monitoring and Decision Support System. In addition, we present infrastructure requirements that will enable a datacenter to be part of the ENEDI ecosystem. For the Monitoring System, we introduce the internal architecture of the components and describe the interactions with the monitored equipment, servers and user applications. We conclude this section with the Decision Support System which is presented through a comprehensive workflow.

A. *ENEDI Monitoring System*

Monitoring holds a prominent position as an individual self-governing system in the Platform Layer. It is responsible for the continuous monitoring process and supervision of the monitoring lifecycle and adheres to the architecture illustrated in Figure 2

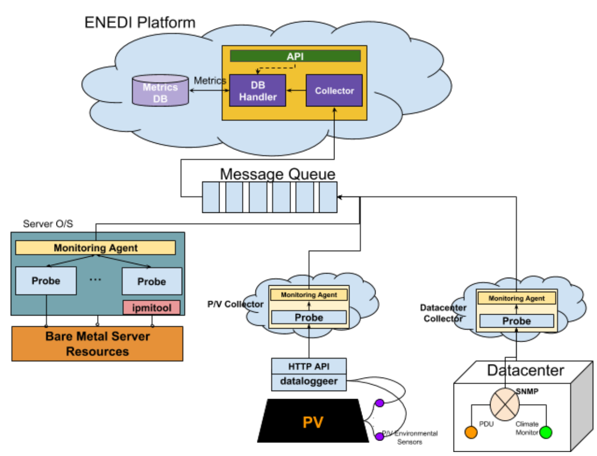

Fig. 2: ENEDI Monitoring System Architecture

The ENEDI Monitoring System adheres to the microservice paradigm and adopts the producer-consumer communication. This design pattern provides interoperable, scalable and real-time monitoring for extracting IoT sensors, operating system and application performance indicators from deployed cloud applications. The ENEDI Monitoring System runs in a non-intrusive and transparent manner to the underlying infrastructure and user applications as neither the metric collection process nor metric distribution and storage are dependent on the underlying platform.

The Monitoring System consists of the following:

- **Monitoring Agents** are processes responsible for coordinating and managing the metric collection process on the respective infrastructure entity (server, IoT sensors, etc.), which includes aggregation and dissemination of monitoring data to the Monitoring Service over a Message Queue. Figure 3 depicts the complete list of sensors that are monitored by ENEDI Monitoring System Additional functionality of a Monitoring Agent includes adapting the periodicity of both the metric collection and dissemination based on the current evolution of the metric data stream.
- **Monitoring Probes** are the actual metric collectors. They feature both a push and pull metric delivery mechanism with Monitoring Agents benefiting from the push mechanism to avoid the overhead of constantly checking for metric updates.
- The **Collector** is responsible for collecting and prometrics from active Monitoring Agents. The communication between Monitoring Agents and the Collector is facilitated by a PUB/SUB communication protocol which reduces the related network communication overhead. Moreover, the component forwards metrics to the Database Handler to be persisted in the ENEDI's Metrics Database.
- The **Database Handler** is the responsible component for persisting monitoring metrics to the ENEDI's distributed, scalable Metrics Database with a high-performance indexing scheme.
- The **Monitoring API** is responsible for providing access to monitoring data stored in the Metrics Database.

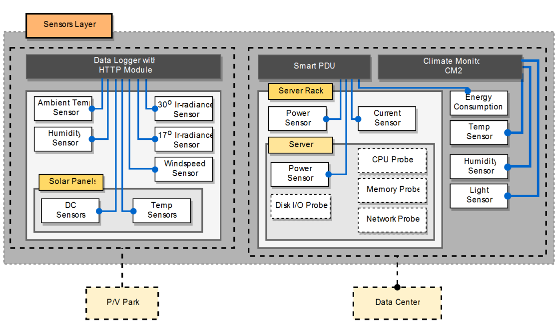

Fig. 3: ENEDI Monitored Sensors and System Probes

B. Decision Support Flow

Figure 4 depicts the decision support flow of the ENEDI Platform. The Data Processing component uses the Monitoring API to retrieve stored metrics from the Monitoring System and perform runtime combilation of high-level analytic insights from the collected monitoring data based on energy policies defined by users. After assessing the validity and correctness of the received metric rules, real-time analytic insights are then served to a Rule Management System that implements pattern-matching algorithms to determine which of the system's rule will fire based on its data store. The data store consists of three distinct data repositories; (a) The Production Memory stores user-defined policies and energy consumption and production thresholds from the Policies Editor component, (b) The Working Memory that feeds the system with facts (metrics) and (c) The Agenda, which contains a list of available rules to be fired. Actions fired from the Agenda consist of the following:

- **Load Management**; ENEDI platform is equipped with a Load Management component.The component is responsible to inform administrators to shut-down or start cloud applications on a specific datacenter.
- **Notifications**; Whenever an agenda action is fired (e.g., re-routing of computational workload to a different datacenter) datacenter administrators are notified through various digital(email, slack, dashboard notification) and analogous(sms, telephone, etc.) means according to the preference of the admin.
- **Logs** Responding actions and triggering events registered in the logs database for future reference.

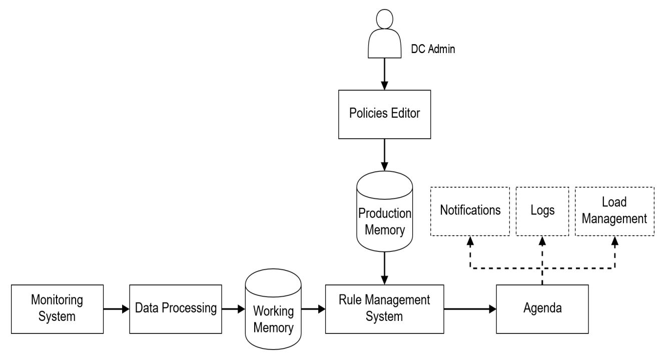

Fig. 4: ENEDI Decision Support Flow

C. Cloud Infrastructure and Virtualization

All datacenter servers participate in a cluster structure which offers cloud resources in the form of VMs, a process which is facilitated by an IaaS middleware management system (Openstack [19] Okeanos [20]). Individually, each server in the cluster is layered as depicted in Figure 5.

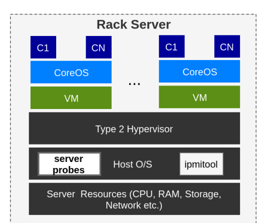

Fig. 5: ENEDI Server Virtualization Layers

Debian-based Linux operating system equiped with Type-2 Hypervisor undertake the task to abstract guest operating systems from the host's by coordinating calls for CPU, memory, disk, network and other resources. Moreover, ENEDI probes that measure server performance are pre-bundled within the host operating system of ENEDI-enabled datacenters along with IPMI which will be used to collect environmental and consumption data from the server sensors. At the top-most layer, ENEDI will offer stateless Docker Container Images as computational workload resources. The main benefits for utilizing container images as computational workload rely on the nature of containers themselves. A Docker Container is a lightweight, stand-alone, executable package that includes everything needed to run a user developed application. Containers exhibit short booting times, portability and vendor-neutrality which enables them to run on-demand on heterogeneous infrastructures in a user-transparent manner. ENEDI's computational workload offering consists of containers within CoreOS VMs with the aforementioned monitoring probes pre-installed. CoreOS has been selected as the VM OS due to its out-of-the-box support for the Docker Engine [21].

D. Photovoltaics and Sensors

Energy consumption and production are of great significance for ENEDI. The primary objective of ENEDI is to minimize operational costs without having to imperil performance of deployed user applications. To achieve this goal, ENEDI will consume energy produced by PV panels in combination with grid energy to power computational, cooling and lighting equipment within datacenters. Moreover, ENEDI employs a network of IoT sensors to retrieve environmental and performance measurements from the PV panels. In Table I we enlist these sensors and their exposed metrics. Due to the analog nature of the IoT sensors, collected metrics propagate to the rest of the system through a data logger. A data logger is a specialized component that can scan a wide variety of measurement, perform any programmed processing on the data and store them in their memory. As depicted in Table I, the measurement variables of temperature (Tm) as well as DC current (IDC), voltage (VDC) and power (Pm) are to be continuously sampled and monitored via the data logger. The latter is not possible to be directly sampled so it will be derived from IDC and VDC. These metrics are to be monitored at a 1-minute resolution and be accumulated as 15-minute averages for the number of selected PV modules.

It is essential for the normal operation of the ENEDI platform, sensor data in data loggers to be accessible to the Monitoring System through an HTTP REST API.

TABLE I: ENEDI List of PV Sensors

PV Sensors List		
Sensor	Metric	Type
Pyranometer	GPOA	Environmental
DC Voltage Sensor	Vdc	Performance
DC Current Sensor	Idc	Performance
Ambient Temp. and Humidity Sensor	^{o}C, %	Environmental
PV Temp. Sensor	^{o}C	Environmental
Wind Speed Sensor	km/H	Environmental
Module Power	P	Performance (Derived)

IV. Conclusion and Future Work

In this paper, we presented the high-level architecture of ENEDI, an integrated system that aims to enable PV-powered distant datacenters to optimize energy consumption, resource utilization and minimize operational costs. ENEDI monitors an extensive network of IoT sensors installed within datacenters and PV panels and combines these measurements with user-application performance measurements and user-provided energy thresholds to determine the location and time-slot of the workload execution.

Currently, ENEDI is at its early development phase. Our efforts focus on the development and configuration of the ENEDI Monitoring System. Specifically, we have developed monitoring agents responsible to retrieve fine-grained insights from user applications, servers, environmental sensors and performance metrics from PVs. Moreover, we have focused on automating the process of creating ENEDI compliant datacenters and the general setup of the infrastructure of the ENEDI platform. In addition, we are using the early version of our Monitoring System to collect performance and environmental metrics from our small-scale datacenter to provide baseline analytics for future comparisons with ENEDI enabled datacenters.

Our next steps aim to finalize the first release of the ENEDI platform and emphasis will be given to the DSS component by implementing its smart rule management system. We also plan to perform real-world experiments by measuring datacenter performance and operational cost improvement of pre and post ENEDI utilization.

Acknowledgment

The ENEDI Project is co-financed by the European Union (ERDF) and national funds of Greece and Cyprus (Cross-Border Cooperation Programme Greece Cyprus 2014 – 2020).

References

[1] Avgerinou, M., Bertoldi, P. and Castellazzi, L., 2017. Trends in datacentre energy consumption under the European code of conduct for datacentre energy efficiency. Energies, 10(10), p.1470.
[2] Google, Renewable energy datacenters, https://www.google.com/about/datacenters/renewable/index.html .
[3] Amazon, AWS and Sustainability, https://aws.amazon.com/about-aws/sustainability .
[4] Microsoft, Transitioning to zero-carbon energy, https://www.microsoft.com/en-us/environment/energy .
[5] Kong, F. and Liu, X., 2015. A survey on green-energy-aware power management for datacenters. ACM Computing Surveys (CSUR), 47(2), p.30.
[6] ENEDI Project Website, https://enedi.eu/ .
[7] Amazon CloudWatch, https://aws.amazon.com/cloudwatch/.
[8] CloudMonix, https://www.cloudmonix.com/aw/.
[9] RackSpace CloudKick, http://www.rackspace.com/cloudkick/.
[10] Massie, M.L., Chun, B.N. and Culler, D.E., 2004. The ganglia distributed monitoring system: design, implementation, and experience. Parallel Computing, 30(7), pp.817-840.
[11] Nagios, http://www.nagios.org/.
[12] Prometheus, https://prometheus.io/ .
[13] Trihinas, D., Pallis, G. and Dikaiakos, M., 2015. Monitoring elastically adaptive multi-cloud services. IEEE Transactions on Cloud Computing, (1), pp.1-1.
[14] Goiri, Í., Beauchea, R., Le, K., Nguyen, T.D., Haque, M.E., Guitart, J., Torres, J. and Bianchini, R., 2011, November. Greenslot: scheduling energy consumption in green datacenters. In High Performance Computing, Networking, Storage and Analysis (SC), 2011 International Conference for (pp. 1-11). IEEE.
[15] Goiri, Í., Le, K., Nguyen, T.D., Guitart, J., Torres, J. and Bianchini, R., 2012, April. GreenHadoop: leveraging green energy in data-processing frameworks. In Proceedings of the 7th ACM european conference on Computer Systems (pp. 57-70). ACM.
[16] Liu, Z., Lin, M., Wierman, A., Low, S.H. and Andrew, L.L., 2011. Geographical load balancing with renewables. ACM SIGMETRICS Performance Evaluation Review, 39(3), pp.62-66.
[17] Dean, J. and Ghemawat, S., 2008. MapReduce: simplified data processing on large clusters. Communications of the ACM, 51(1), pp.107-113.
[18] Trihinas, D., Tryfonos, A., Pallis, G., and Dikaiakos, M., 2018. DevOps as a Service: Pushing the Boundaries of Microservice Adoption. IEEE Internet Computing, 22(3), pp.65-71.
[19] Openstack, https://www.openstack.org/
[20] GRNet Okeanos, https://okeanos.grnet.gr/home/ .
[21] Docker, https://www.docker.com/ .