Consiglio Nazionale delle Ricerche
Istituto di Calcolo e Reti ad Alte Prestazioni

Performance Evaluation of EcoMultiCloud, a Hierarchical Algorithm for Green and Cost-Efficient Workload Management in Geo-distributed Data Centers

Agostino Forestiero [1,2], Demetrio Laganà[1],
Carlo Mastroianni[1,2]

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 712363

RT-ICAR-CS-17-01 **October 2017**

1) Eco4Cloud srl, Piazza Vermicelli, Rende (CS), email: lagana@eco4cloud.com
2) ICAR-CNR, Via P. Bucci 7/11 C, Rende (CS), email: {forestiero,mastroianni}@icar.cnr.it

Consiglio Nazionale delle Ricerche, Istituto di Calcolo e Reti ad Alte Prestazioni (ICAR)
– Sede di Cosenza, Via P. Bucci 7-11C, 87036 Rende, Italy, URL: www.icar.cnr.it
– Sezione di Napoli, Via P. Castellino 111, 80131 Napoli, URL: www.icar.cnr.it
– Sezione di Palermo, Via Ugo La Malfa, 153, 90146 Palermo, URL: www.icar.cnr.it

1 Introduction

In this report, we describe how the algorithms for workload assignment and redistribution of the EcoMultiCloud architecture, presented in [1], are customized for a specific multi-data center environment, i.e., the one owned and managed by the Italian telecommunications company TIM, formerly Telecom Italia.

The scenario considered in this document is composed of three interconnected DCs of the TIM company across Italy. The DCs are located respectively in Rozzano (MI) - DC1, Pomezia (RM) - DC2 and Bari (BA) - DC3. For the sake of simplicity the three DCs are considered equal in architecture, number and capacity of hosts. The presence of three equal data centers permits to abstract the workload distribution from the specific architectures and sizes of the three data centers and use them as a testbed for a deep analysis assessment and for the tuning and optimization of the EcoMultiCloud software. Moreover, this approach allowed us to achieve interesting results without using and publishing confidential data about the TIM hardware facilities.

The rest of the report is structured as follows: Section 2 individuates, in cooperation with TIM administrators, the business and technical goals that are of primary interest for the company, and gives details about how they are computed and evaluated. Then, the section describes how the algorithms for workload assignment and redistribution are customized and tuned to match the selected goals.

Section 3 illustrates the adopted testbed and describes the type of energy produced in the three data centers. All the data centers use a mixture of green and grid energy, but produce different types of green energy: exclusively solar, exclusively wind, or both solar and wind. We also adopted a dynamic approach to compute the Power Usage Effectiveness (PUE) of the data centers, based on the workload and the external temperature.

Section 4 describe the obtained results. Specifically, Section 4.1 shows the values of the main performance indices, i.e., the energy cost, the carbon emissions, the amount of energy, also differentiated by source, the workload distribution, and the number of migrations. The results prove the the EcoMultiCloud algorithm is very effective and also very flexible, in that it is possible to tune the weights of the different business goals, depending on the choices of the administrators. In Section 4.2 we report the feedback received by TIM administrators. In particular, they asked us to better focus on the amount of available bandwidth, as it can vary depending on specific requirements, and on the number of migrations, which should be minimized in order to reduce the impact on the quality of service. Following these advices, we refined the policy for the selection of the VMs to migrate between data centers. Finally, in Section 4.3 we present the results when varying the available bandwidth for remote migrations, and show that the algorithm is effective even when the bandwidth is limited. Furthermore, we show that the refined policy for VM selection allows the number of migrations to be reduced.

Finally, Section 5 concludes the document.

2 Refinement of the Algorithms for Workload Distribution

The EcoMultiCloud solution has been designed and implemented upon a hierarchical architecture, which allows to obtain the following benefits: (i) better scalability, since the problem is decomposed in two different layers, and the size is notably reduced; (ii) better ability to adapt to changing conditions; (iii) improved autonomy of the sites, which are not required to share the same strategies and algorithms.

The software is composed of two layers: (i) the *upper layer* is used to exchange information among the different sites and drive the distribution of VMs among the DCs and (ii) the *lower layer* is used to allocate the workload within single DCs. The reference scenario is depicted in Figure 1, which shows the upper and lower layer for two interconnected DCs, as well as the main involved components. At each DC, a data center manager (DCM) runs the algorithms of the upper layer, while the local manager (LM) performs the functionalities of the lower layer.

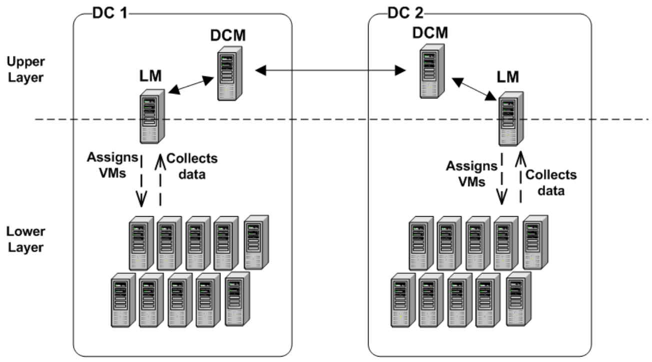

Fig. 1. EcoMultiCloud scenario: upper and lower layer of two interconnected data centers.

The workload management is performed using a set of algorithms, among which: (i) the *assignment algorithm* that determines the data center and the server on which a new VM should be allocated; (ii) the *DCSA algorithm* (Data Center Selection Algorithm) that determines when a workload migration should be triggered and determines the source and the target data center of the migration and (iii) the *VMSA algorithm* (Virtual Machine Selection Algorithm) that determines which type of virtual machines should be migrated.

All the three algorithms make use of an assignment function, evaluated at each data center. This function combines and weighs a number of different terms that are related to different business goals that must be accomplished. Even if the general form of this function is designed to be applicable to each scenario, it is necessary to customize it in accordance with the specific set of goals defined for each specific environment.

To simplify reading, in Table 1 the notation used throughout the document is reported.

Table 1. Symbols and notation.

N_{DC}	Number of DCs
f^i_{assign}	Assignment function for DC i
W_s	Power consumed by server s
F_d $\widehat{F}_d$ ($\widehat{F}_{max}$)	Carbon emission rate for DC d Carbon emission rate $\cdot$ PUE for DC d (max c.e. among the DCs)
U_d (U_{max})	Utilization of DC d (max utilization among the DCs)
P_d (P_{max})	Price of energy at DC d (max price of energy among the DCs)
α, β, γ	Weights of the terms considered in the f_{assign}
PUE_d	Power Usage Effectiveness for DC d
Φ_d	Carbon emissions for DC d
C^E_d	Energy cost for DC d
C^Φ_d	Cost of carbon emissions for DC d
Λ	Total workload

2.1 Identification of Business and Technical Goals

With the help of the administrators of the TIM data center environment, we individuated the business and technical goals that should be addressed through an appropriate tuning of the EcoMultiCloud algorithms.

The most relevant business and technical goals for the TIM environment are:

Reduction of Energy Consumption The minimization of energy consumption is a primary goal for TIM, as this allows to maximize the efficiency of the infrastructure and to improve the utilization of both the computational resources – i.e., CPU, RAM, disk – and the physical resources of the data center, specifically, the power distribution facilities and the cooling components. Moreover, the maximization of green energy usage is considered an important objective, as highlighted in the 2016 TIM sustainability report[1].

In this context, EcoMultiCloud can be very effective, both by consolidating the workload on the minimum number of physical machines, and by redistributing the workload on the most efficient servers and on the different data centers.

The consumed energy can be computed multiplying the power consumed by the servers by the PUE of the corresponding data centers, in order to consider the energy consumed by the cooling and power distribution components. This aspect is illustrated in Figure 2, which shows the energy chain in a data center.

The total energy consumed in the data centers in a given time interval is computed as:

$$E = \sum_{d,s} \int W_s(t)dt \cdot PUE_d(t) \tag{1}$$

[1] http://www.telecomitalia.com/tit/it/sustainability/sustainability-report/sustainability-reports.html

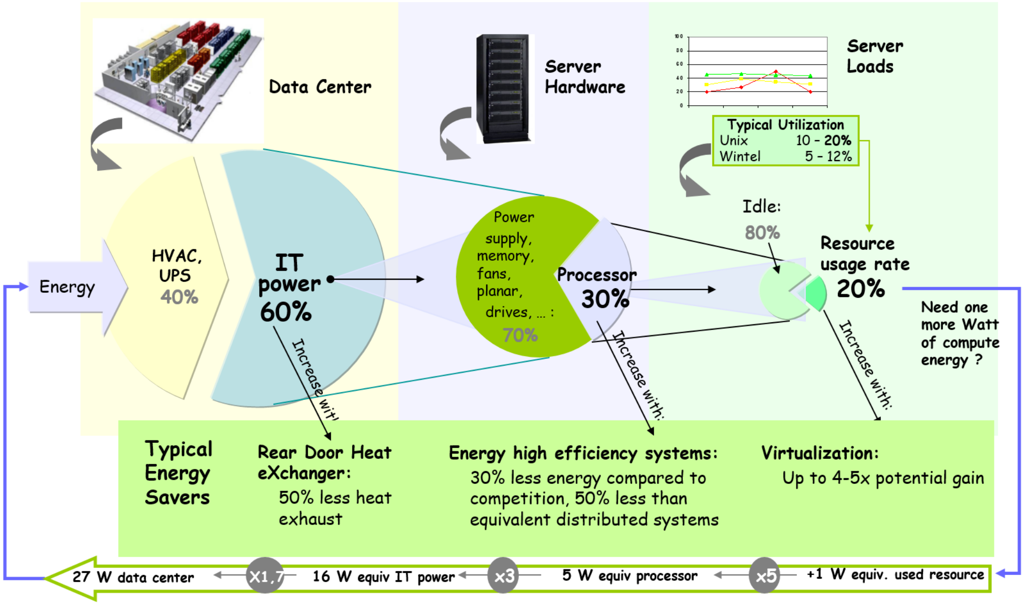

Fig. 2. Energy chain in a data center.

In this expression, the energy consumed by a single server s is computed by integrating the consumed power, $W_s(t)$, over the considered time interval. Then, the total energy is obtained by summing up the values of consumed energy for each server s and each data center d. It is noticed that the value of the PUE, $PUE_d(t)$, depends on the data center and is a function of the time t. Indeed, it will be shown in Section 3.2 that the PUE is not constant but depends on several parameters, among which the external temperature and the workload.

Expression (1) can be easily customized to compute the amount of green/brown energy consumed in the data center.

We recall here that the power $W_s(t)$ consumed by a server mostly depends on the CPU utilization, as described in [1].

Reduction of Carbon Emissions The impact of operational carbon usage is emerging as extremely important in the design and operation of current and future data centers. Specifically, this business objective is considered primary for TIM, as this company is willing to improve its "green" image and prove the efficiency of its efforts for a better protection of the environment, as described in the annual report mentioned in the previous section.

The two most indices used to measure this phenomenon are the CUE (Carbon Usage Effectiveness) of the data center, and the value of total carbon emissions. Here we focus on the second index, as our intention is not to improve the carbon effectiveness of the data center, which would involve infrastructure modifications, but to minimize the carbon emissions in the existing data center sites.

The carbon emissions are determined from the geographical location and from the actual mix of energy delivered to the site, i.e., the electricity generated from varying CO2-intensive plants-coal or gas and from hydro or wind sources. Moreover, the energy can be extracted from the public power grid or from private plants.

The emissions can be derived through the *carbon footprint rate*, which is defined as the amount of CO2 kilograms emitted per energy unit. The computation of the carbon footprint rate covers the operations of the servers and the data centers, but it does not cover the full environmental burden of the life-cycle, including, for example, the carbon generated in the manufacturing of the IT equipment and its subsequent shipping to the data center.

The carbon footprint rate is denoted as F in this document and is defined as:

$$F = \frac{CO_2 \ emitted \ (kg)}{energy \ unit \ (kWh)} \tag{2}$$

In our environment, each data center uses a fraction of renewable energy and a fraction of energy taken from the power grid, or "brown" energy, and the amounts of such fractions depend on the current availability of green energy. The two types of energy are distributed among the servers, so as to use green energy to feed as many servers as possible, while the remaining servers use brown energy.

The total amount of carbon emissions in a given time interval is computed as:

$$\Phi = \sum_{d,s} \int W_s(t)dt \cdot PUE_d(t) \cdot F_d \tag{3}$$

In this expression, the carbon emission rate is zero for servers that use green energy, while it is equal to F_d for servers that use brown energy in the data center d.

Reduction of Monetary Cost The overall monetary cost is the sum of the cost related to the energy consumption and the cost related to the carbon emissions.

When computing the cost related to the energy consumption, the source of the energy must be considered. Indeed, the price of the energy depends if the energy is taken from the power grid or if it is produced locally by green energy sources, such as photovoltaic panels or wind farms. In all cases, i.e., for each of the considered energy sources, the price changes with time. This variability depends on the energy market (for example, the price of energy is lower during the night and in the weekends) and on the availability of green energy, which is unpredictable to a large extent.

For simplicity, we set the price of green energy, produced locally, to zero, without considering the depreciation of the Capex costs sustained to build the energy plants. Therefore, the energy cost due to energy consumption is computed as:

$$C^E = \sum_{d,s} \int W_s(t)dt \cdot PUE_d(t) \cdot P_d(t) \tag{4}$$

In this expression, the price of the energy used at the data center d at time t, $P_d(t)$, is valued only for the servers that use energy taken from the power grid. The cost of energy for the data center d is denoted as C_d^E.

The cost of carbon emissions is related to the carbon tax T_C, i.e., the tax that is applied to the emissions of CO2. Therefore, it can be computed as the amount of carbon emissions multiplied by the tax:

$$C^{\Phi} = \Phi \cdot T_C \tag{5}$$

The cost of carbon emissions for the data center d is denoted as C_d^{Φ}.

In conclusion, the total cost for the entire multi-data center environment is $C = C^E + C^{\Phi}$, while the cost for the data center d is $C_d = C_d^E + C_d^{\Phi}$.

Load Balance The business objectives discussed so far must take into account the fact that a proper load balance must be maintained among the different data centers, i.e., even if a data center, at a given time, is more convenient than the others, the workload cannot be distributed so that this data center gets fully loaded while other data centers are left almost unloaded.

A proper balance of load is required for several reasons, among which:

– the convenience of a data center changes with time, therefore the assignment of a too large fraction of load to a single data center can be a bad choice;
– a portion of workload is related to applications and virtual machines that must run on a specific data center and cannot be migrated;
– a proper balance of load also means a proper balance of the maintenance and administrative tasks that must be carried out by the company employees;
– a proper load balance minimizes the risks of overload, which can lead to a low quality of service offered to users.

In our environment, the bottleneck hardware resource is the RAM memory, so the objective is to limit the imbalance among the RAM utilization at the different DCs. The RAM utilization at the data center d is denoted as U_d and is defined as the ration between the amount of utilized RAM and the maximum allowed utilization.

Technical Goals Beyond the business objectives, it is also important to match a set of technical goals and constraints, which concern two main aspects:

– Service interruptions and delays perceived by the users (QoS degradation)
– Impact on resources utilization

The active action taken by the EcoMultiCloud software to improve such aspects is to perform long distance migrations among data centers. We performed a full analysis of the migrations executed on our test environment, with the aim to evaluate the effect on the QoS degradation perceived by the user and the impact on resources utilization, like the CPU and RAM overhead and the network bandwidth utilization.

The parameters that are used to characterize the technical goals are:

- Migration Duration
- Packet Loss
- Cumulative downtime
- Latency
- Needed bandwidth
- Host overhead (CPU and RAM)

During our tests, we noticed that the migration duration has no effect on the QoS degradation perceived by the users. Indeed, when a VM migrates from a site to another in real-time, the virtualization infrastructure builds a copy of the VM on the destination host. In the meanwhile all the requests sent to the VM are managed by the source copy, so the critical moment of this process is near the end of the migration, when there is the need to switch all the remote requests from the source to the destination. On the other hand, the migration duration has a big impact on resources utilization, as it is strictly correlated to the VM size, in terms of RAM and disk, and to the bandwidth needed to transfer the VM data. Packet loss, cumulative downtime and latency are also strictly connected to the experienced interaction of the user with the VM and are independent from the type of VM. The bandwidth and host overhead are also important because they help to measure the additional need for hardware resources during the VM migrations.

In conclusion, the challenge is to reach the business goal, while respecting the technical constraints, minimizing the use of hardware resources and reducing the impact perceived by end users.

2.2 Customization of the Algorithm for Workload Placement and Migration

As mentioned earlier in this section, the algorithms for workload assignment and migration make use of an assignment function f_{assign}, evaluated at each data center. This function combines and weighs a number of different terms that are related to different business goals that must be accomplished. In our scenario we decided, in cooperation with TIM managers and data center administrators, to customize the assignment function by considering three business goals discussed so far: the reduction of carbon emissions, the reduction of monetary costs and the workload balance among the data centers. The reduction of energy consumption is not specifically included in the function, but it is strictly related both to the carbon emissions and to the monetary cost. This relationship is evident in Expressions (3) and (4).

The technical indices, on the other hand, are not included in the function f_{assign}, because they are best seen as technical constraints, especially for the migration procedures. In practice, this means that the workload is migrated in accordance to the decisions of the migration algorithms, but during the remote migrations of Virtual Machines, the technical indices are continuously monitored. When it occurs that one of these indexes does not respect the corresponding constraint (e.g., the downtime experienced by users during the migrations is excessive), the migration procedure is suspended.

The general expressions of the algorithms were presented in the scientific paper [1]. Here we describe how the assignment function is customized in order to match the three mentioned business goals. The assignment function associates to each DC a value that represents the cost to run some workload in that DC, low values correspond to low overall cost of the DC. The strategy, then, is to assign a VM to the DC with the lowest value of the function. Here the the objectives are the balance of load, the minimization of carbon emissions and the minimization of costs related to energy, so the assignment function f_{assign}^{d}, for each DC d, is defined as follows:

$$f_{assign}^{d} = \alpha \cdot \frac{\widehat{F}_d}{\widehat{F}_{max}} + \beta \cdot \frac{P_d}{P_{max}} + \gamma \cdot \frac{U_d}{U_{max}} \tag{6}$$

where the coefficients α, β and γ are positive and $\alpha + \beta + \gamma = 1$.

The terms $\widehat{F}_d$, P_d and U_d are related, respectively, to carbon emissions, energy costs and overall utilization. Specifically, the term $\widehat{F}_d$ is equal to the carbon emission rate F_d multiplied by the value of PUE_d. The terms are normalized with respect to the maximum values communicated by DCs. The proper balance among the three goals is determined by the strategic decisions taken by the system administrator, and is tuned through the values of the parameters α, β and γ. After computing the values of f_{assign} for each DC, the VM is assigned to the data center having the lowest value. Once consigned to the target DC, the VM is allocated to a physical host using the local assignment algorithm EcoCloud [3].

To compute the assignment function, each data center d periodically (i.e., every hour) communicates to the central manager, and to the other data centers, some very simple pieces of data. In the examined case, relevant information is:

1. the best available carbon footprint rate of a local server. It is equal to zero if there is some available and not exploited green energy, otherwise it is equal to the rate F_d, measured in Tons/MWh[2];
2. the current value of the PUE;

[2] Note that, when assigning a VM, the target DC should be chosen so as to minimize the incremental increase of the carbon footprint. To this aim, a DCM does not need to know the carbon footprint rate of all the servers of remote sites: it only needs to know, per each site, the minimum rate among the servers that are available to host the VM.

3. the utilization of the bottleneck resource, U_d, in our case the RAM memory;
4. the current value of the price of energy P_d;

In conclusion, each data center transmits to the other data centers the following vector of values, which corresponds to the *state* of the DC:

$$s_i = \{F_d, PUE_d, U_d, P_d\} \tag{7}$$

Figure 3 reports the pseudo-code used by a data center DCM (Data Center Manager) to choose the target data center, among the N_{DC} data centers of the system, for a VM originated locally. First, the DCM requests the values of F_d, U_d, PUE_d and P_d to all the other data centers. Then, it computes the maximum values of the terms, for the normalization, and computes expression (6) for any data center that has some spare capacity, i.e., for which the utilization of the bottleneck resource has not exceeded a given threshold U_{T_i}. Finally, the VM is assigned to the DC that has the lowest value of (6). Once consigned to the target DC, the VM is allocated to a physical host using the local assignment algorithm.

This algorithm tends to distribute the workload so as to make all the data centers equally convenient, i.e., the values of the assignment functions at the different data centers tend to be comparable. The procedure for workload migration is triggered when the values of f_{assign} for two data centers differ by more than a predefined threshold, e.g., 3%

function AssignmentAlgorithm(α, β, γ)
  while VM arrives
    for each remote datacenter DC_d
      Request values of F_d, PUE_d, U_d, C_d
      $\widehat{F}_d = F_d \cdot PUE_d$
    end for
    $\widehat{F}_{max} = Max\{\widehat{F}_d |\ d = 1 \cdots N_{DC}\}$
    $U_{max} = Max\{U_d |\ d = 1 \cdots N_{DC}\}$
    $P_{max} = Max\{P_d |\ d = 1 \cdots N_{DC}\}$
    for each DC_d : DC_d is not full, that is, $U_d < U_{T_d}$
      $f^i_{assign} = \alpha \cdot \frac{\widehat{F}_d}{\widehat{F}_{max}} + \beta \cdot \frac{C_d}{C_{max}} + \gamma \cdot \frac{U_d}{U_{max}}$
    end for
    $DC_{target} = DC_j$ such that $f^j_{assign} = min\{f^i_{assign} |\ i = 1 \cdots N_{DC}\}$
    Assign VM to DC_{target}
  end while
end function

Fig. 3. The EcoMultiCloud assignment algorithm, executed by the DCM of each data center.

3 Testbed Description

This section illustrates the description of the environment used to perform the experiments and evaluate the performances of the system. In section 3.1 we provide a complete description of the environment and data parameters used in the EcoMultiCloud simulator, in section 3.2 we describe the different energy sources for each data center.

3.1 Environment Definition and Characterization

In this section we describe the parameters used to asses the EcoMultiCloud scenario. The environment is emulated using an event-based Java simulator. In this work, the software has been enriched, adding the new features to support the new requirements. The considered scenario is composed of three interconnected DCs of the TIM company across Italy. The DCs are located respectively in Rozzano (MI) - DC1, Pomezia (RM) - DC2 and Bari (BA) - DC3. For the sake of simplicity the three DCs are considered equal in architecture, number and capacity of hosts. This choice allows us to better understand the behavior of the software, specifically when the availability of green energy and the convenience of the single data centers changes with time. The presence of three equal data centers permits to abstract the workload distribution from the specific architectures and sizes of the three data centers.

The data center asset is composed of 28 hosts, virtualized with the platform VMware vSphere. All the servers are equipped with processor Xeon with a clock frequency varying from 2.4 to 2.9 GHz, a number of cores in the range 16-24 and a number of physical sockets of 4. All the servers are equipped with a total amount of physical RAM varying from 64GB to 96GB, also have network adapters with a bandwidth of 1 Gbps. The servers hosts 414 VMs which are assigned a number of virtual cores varying between 1 and 8 and an amount of assigned RAM varying between 256 MB and 16 GB. This data are collected from the data center located in Pomezia (DC2). Also data concerning the workload is gathered from the data center DC2, and replicated in the other two data centers, so as to set the DC load to Λ, see Table 1.

The detailed data about the TIM servers, the corresponding value of hardware resources, and the data about the real workload, are used in the experiments but are not reported here because they are confidential data. Moreover, as explained above, we simplified the scenario, by using the same number of hosts and the same computational infrastructure in the three data centers, in order to better assess the performance and effectiveness of the adopted algorithms and procedures. In addition, some data regarding the real TIM environment, for example the values of the PUE, the availability of renewable energy sources and data about the energy production cannot be used in this work because they are confidential or because they are not available. Therefore we decided to use data taken from the literature. This data is provided in the next section.

3.2 Energy Sources

In our environment, three energy sources are considered:

- Solar energy
- Wind energy
- Grid energy

In our test, Solar and Wind energy are considered as green energy, with a carbon footprint rate and energy cost equal to zero. The grid energy has a gross monetary cost of 0.1138 €/kWh([3]), while the carbon footprint rate is set to 0.350 Tons/MWh, as in [2]. Each DC is connected to the grid and has a nominal energy power of 50 kW. This data is reported in Table (2).

Table 2. Price of grid energy, expressed as €/kWh, carbon footprint rate of grid energy, expressed in Tons/MWh and nominal energy power, expressed in kW for the 3 DCs.

	DC_1	DC_2	DC_3
Energy Price (€/kWh)	0.1138	0.1138	0.1138
Carbon Footprint Rate (Tons/MWh)	0.350	0.350	0.350
Nominal Energy Power (kW)	50	50	50

Renewable energy sources provide a portion of the electricity needs of the DCs and are available as follows:

- DC1 is powered by the grid and by solar energy.
- DC2 is powered by the grid and by wind energy.
- DC3 is powered by the grid and by solar and wind energy.

The renewable energy sources have the priority on the grid source, meaning that each data center first exploits all the available solar and wind energy, and uses the grid energy only when the green energy is not sufficient. The availability of renewable energy depends of the energy production on the different sites. We use a distribution of green energy production for the three data centers as follows: Figure 4 reports the typical trend of solar energy production at the DC 1; Figure 5 reports the trend of wind energy production at the DC 2; finally, Figure 6 reports the cumulative production at the DC 3 of solar and wind energy[4]. In these figures, the value of the y axis represents the fraction of produced green energy, i.e., the ratio between the green energy produced at each hour of the day and the nominal energy needed to power the data center at full utilization. The nominal energy for each single data center is set to 50 kW. It is important to notice that the

[3] http://www.autorita.energia.it

[4] these trends are in part derived from https://blogs.scientificamerican.com/solar-at-home/a-solar-detective-story-explaining-how-power-output-varies-hour-by-hour/

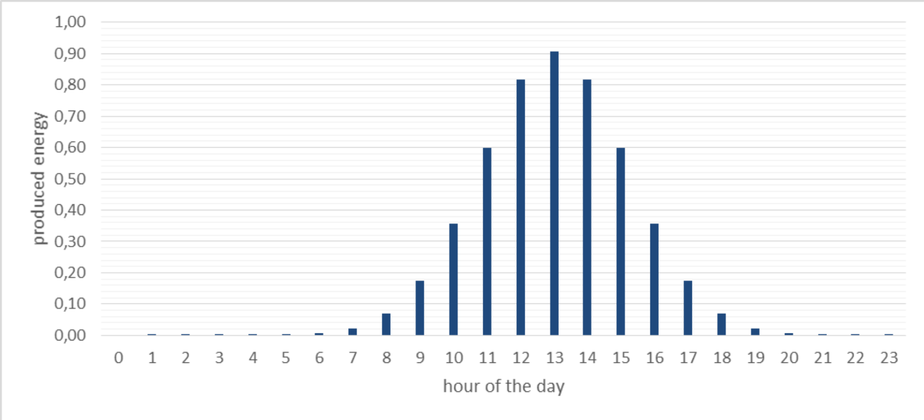

Fig. 4. Solar energy production in Rozzano Data Center (DC1).

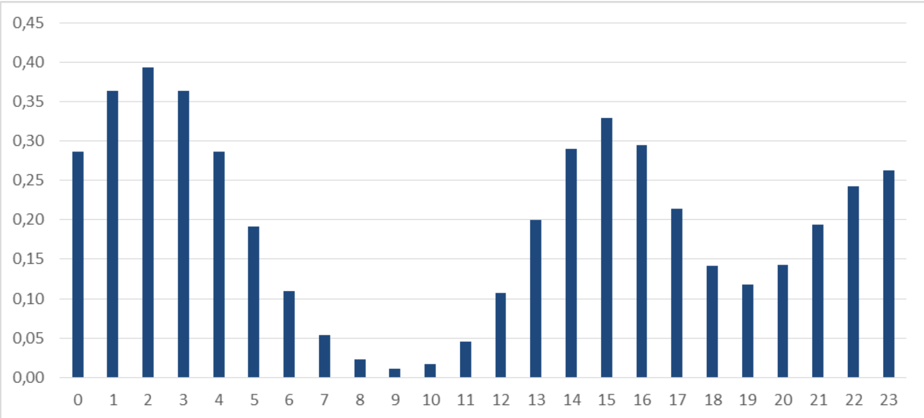

Fig. 5. Wind energy production in Pomezia Data Center (DC2).

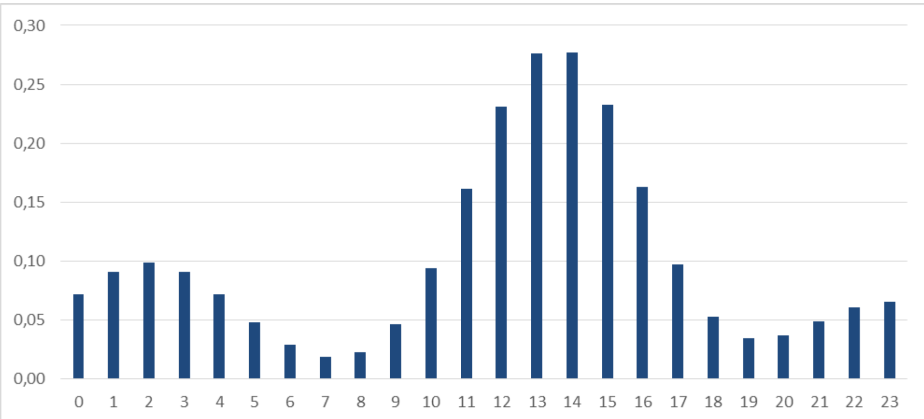

Fig. 6. Solar and wind energy production in Bari Data Center (DC3).

overall green energy produced by DC 1 and DC 2 in 24 hours is comparable, while the green energy produced by DC 3 is about 50% of the other two DCs.

We adopt a dynamic model for the PUE: the PUE is determined on the basis of the temperature and the load of the data center. We adopt the mathematical model defined and described in [4]. Table 3 shows the temperatures, expressed in Celsius degrees, of each DC in a typical day of July.

Table 3. Temperature, expressed in Celsius degree, for the three DCs. The table shows the time, expressed in GMT+1, corresponding to the temperature of each DC in a typical day of July.

Time (GMT+1)	DC_1	DC_2	DC_3
0:00	24	23	26
1:00	24	23	25
2:00	23	22	24
3:00	23	21	23
4:00	23	21	23
5:00	23	21	23
6:00	24	23	25
7:00	25	25	28
8:00	26	28	32
9:00	26	29	34
10:00	27	30	35
11:00	28	30	36
12:00	29	30	36
13:00	30	31	37
14:00	30	31	36
15:00	31	31	36
16:00	31	31	35
17:00	30	30	34
18:00	29	30	33
19:00	28	29	32
20:00	28	27	31
21:00	26	26	29
22:00	25	25	28
23:00	25	24	27

4 Performance Evaluation

This Section is devoted to the performance evaluation of EcoMultiCloud for the testbed environment described previously. In subsection 4.1, we report the value of the performance indices when using three different set of values for the parameters α, β and γ, defined and used in Expression (6). In these experiments, the value of the available bandwidth is set to 2 Gbps. Then, in subsection 4.2, we report some interesting feedback received from the data center administrators of the TIM data centers. Such feedback are aimed to tune the amount of available bandwidth and the number of migrations. To match such advices, we devised some further optimization of the EcoMultiCloud algorithm. Finally, in subsection 4.3, we report the results obtained after the algorithm optimization.

4.1 Business and Technical Results

In the following, we will show the values of the main performance indices, related to business and technical goals, for a 24 hour interval. In order to show steady results, i.e., after the end of the transient phase that follows the activation of the software, we consider the values measured in the third day of operation.

The reported indices, per hour, are:

- the carbon emissions for the three data centers, and the overall values
- the energy cost for the three data centers, and the overall values
- the consumed energy for the three data centers, and the overall values
- the consumed energy, differentiated for energy source (solar, wind and grid)
- the distribution of the workload on the three data centers
- the number of migrations for the three data centers

In the following subsections, we report the values of the mentioned indices for three set of values of the parameters α, β and γ, defined and used in Expression (6).

The adopted parameter values are:

- **Setting 1.** α=0.33, β=0.33 and γ=0.34. With this setting, the three individuated business goals, i.e., the carbon emissions, the cost of energy and the load balance, are given the same relevance;
- **Setting 2.** α=0.5, β=0.5 and γ=0. With this setting, no relevance is given to the load balance, so the objective is to minimize the carbon emissions and the costs.
- **Setting 3.** α=0.4, β=0.4 and γ=0.2. With this intermediate setting, more relevance is given to the minimization of carbon emissions and costs, but the load balance is also taken into consideration, though with less relevance.

Results with Setting 1

Setting 1 corresponds to the values α=0.33, β=0.33 and γ=0.34 in Expression (6).

Figure 7 reports the distribution of the workload on the three data centers on the third day of operation after the activation of the EcoMultiCloud software. The workload corresponds to the fraction of utilized RAM with respect to the overall RAM capacity of each data center. Indeed, the RAM memory is the bottleneck hardware resource in the three data centers. The figure shows that the utilization of the data centers varies with the hour of the day, depending on the availability of the green energy produced on the three data centers. Specifically, we notice that in the central hours of the day, a significant portion of the workload is moved from DC2 to DC1, since the latter produces a larger amount of green (solar) energy. Moreover, we notice that the workload assigned to DC3 is lower with respect to the other two DCs, because this data center produces, in the 24 hours, about half the green energy of the other two DCs.

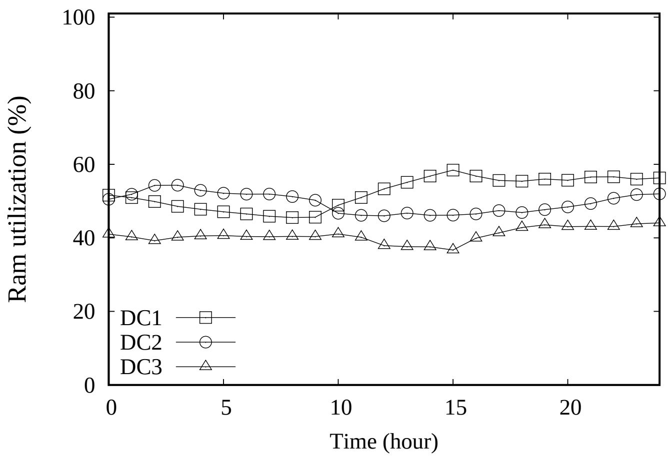

Fig. 7. RAM utilization on the three data centers.

Figure 8 reports the values of consumed energy (considering both green and grid energy) for the three data centers, and the overall value, for the third day of operation. Since the three data centers are equipped with the same hardware infrastructure, we can notice that the trend of consumed energy on the three data centers partly follows the respective resource utilization, reported before. For example, we notices a peak of consumed energy on DC1 in the central hours of the day.

In Figure 9, we report the value of overall consumed energy differentiated by energy source, i.e., solar, wind and grid. We see that the consumption of green energy follows the trend of energy production, as reported in Figures 4, 5 and 6. It is also noticed that the grid energy is not always fully utilized: at some hours of the day, the use of grid energy is very low, while in other moments it is heavily used to complement the insufficient production of green energy.

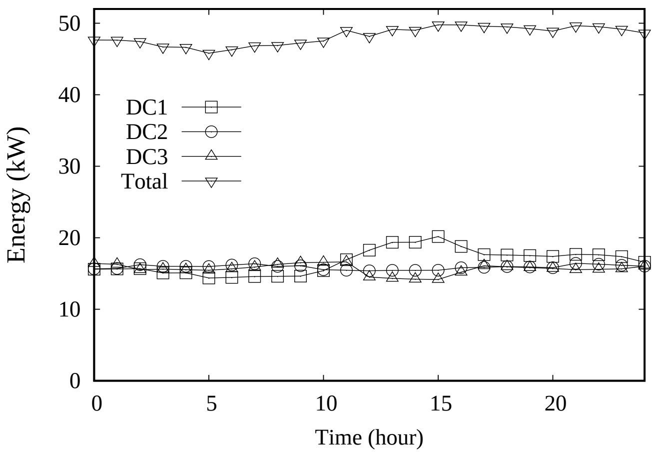

Fig. 8. Consumed energy for the three data centers, and overall value.

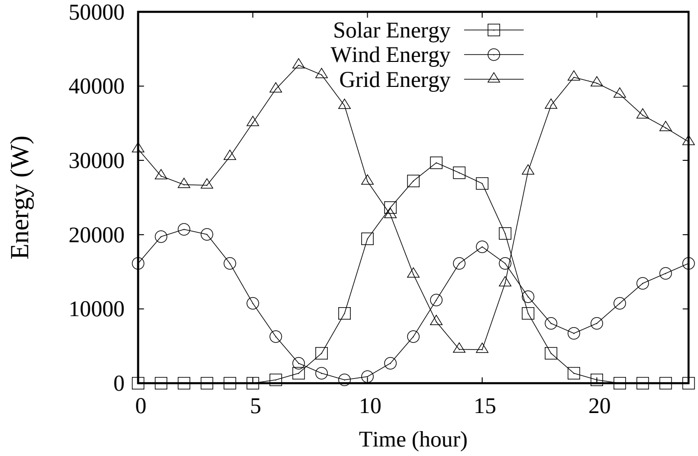

Fig. 9. Energy consumed on the data centers, differentiated for energy source.

Figure 10 shows the value of carbon emissions for the three data centers, and the overall values. We can see that the amount of carbon emissions follows the usage of grid energy, reported in the previous figure. For example, we can see that the carbon emissions are very low in the central hours of the day, between 12:00 and 16:00, when the solar and wind energy are sufficient to cover the energy needs of the data centers.

In Figure 11 we see that the energy cost follows the same trend as the carbon emissions. Indeed, both indices are related to the amount of grid energy. We recall here, indeed, that the cost of green energy is set to zero, because it is produced locally, and the carbon emissions of green plants are also equal to zero.

Finally, in Figure 12 we report the number of migrations performed to optimize the workload balancing with the adopted values of the parameters of the assignment function (6). The number of migrations is limited by the available bandwidth, which is set to 2 Gbps. We can notice the large number of migrations involving DC1 and DC2 between 11:00 and 17:00, as anticipated in the comment to Figure 7.

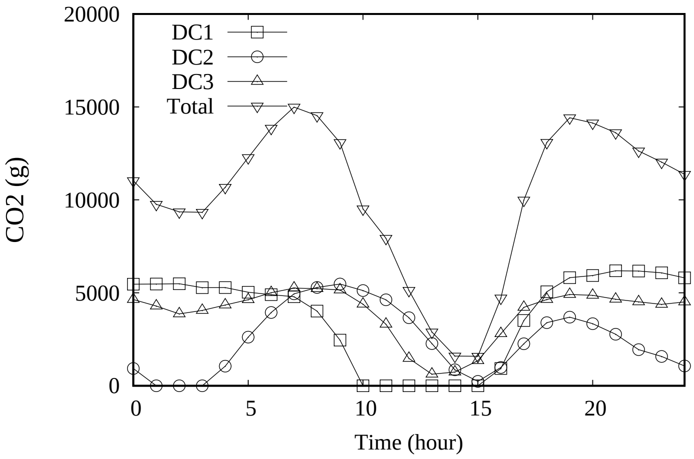

Fig. 10. Carbon emissions for the three data centers, and overall value.

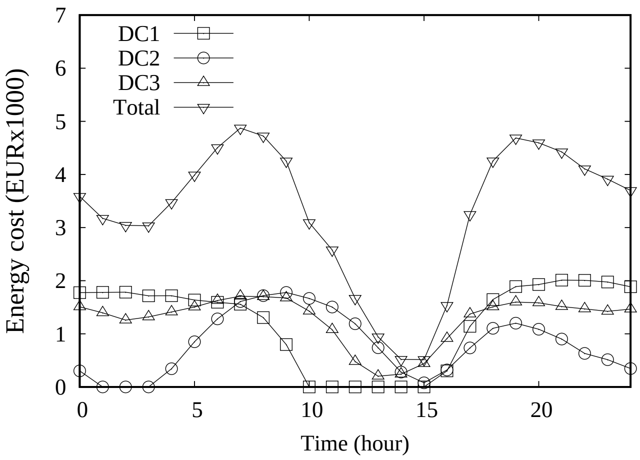

Fig. 11. Energy cost for the three data centers, and overall value.

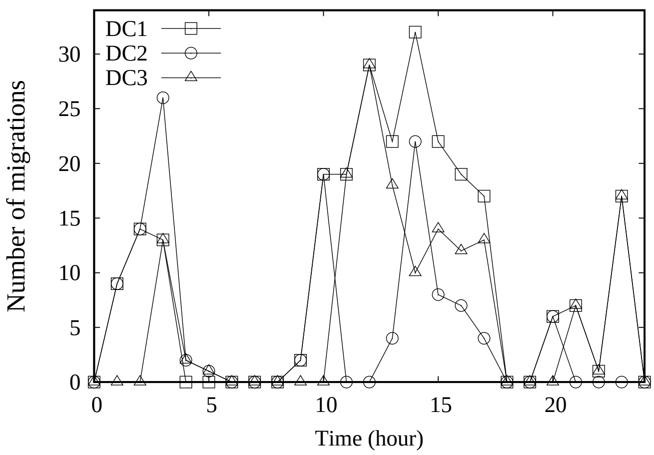

Fig. 12. Number of migrations on the three data centers.

Results with Setting 2

Setting 2 corresponds to the values α=0.5, β=0.5 and γ=0 in Expression (6). With this setting, no relevance is given to the load balance, therefore the data centers that are more efficient in terms of carbon emissions and energy cost are allowed to take any portion of the load by means of VM migrations.

Indeed, Figure 13 shows that the utilization of the data center 3 is much lower than the other data centers. The reason is that data center 3 has less availability of green energy during the 24 hours, therefore the value of the assignment function (6) is lower than the value of the other two data centers. A similar trend is noticed in Figure 14, which shows the hourly energy consumed by the three data centers.

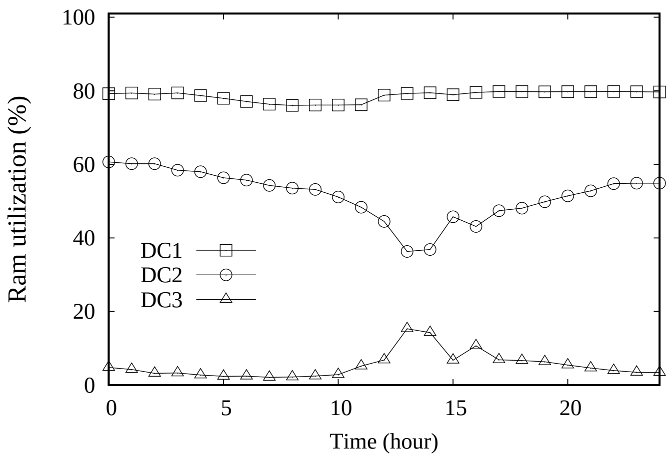

Fig. 13. RAM utilization on the three data centers.

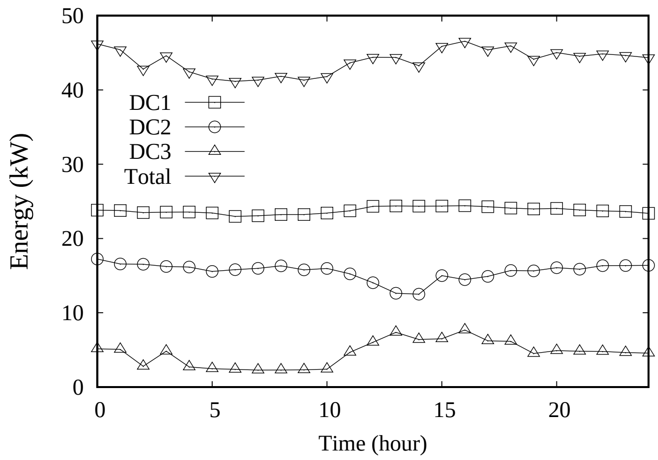

Fig. 14. Consumed energy for the three data centers, and overall value.

Figure 15 shows the amount of consumed energy, differentiated by the energy source. The trend is similar to the one observed in Figure 9, since it is related to the amount of produced energy at the three data centers. However, we can notice that the amount of grid energy is lower: the reason is that no relevance is given to the load balancing, therefore it is possible to move a larger portion of the load to the data centers that are powered by green energy.

The trend of the carbon emissions and of the energy cost, respectively shown in Figures 16 and 17, follow the value of the consumed gird energy, reported in Figure 15. We can notice that the overall energy cost is reduced with respect to Setting 1 (see Figure 11), because a larger weight is given to the energy cost. In the 24 hours, the energy cost is about € 78800 with Setting 1, while it is € 66300 with Setting 2.

Finally, Figure 18 reports the number of performed migrations per hour.

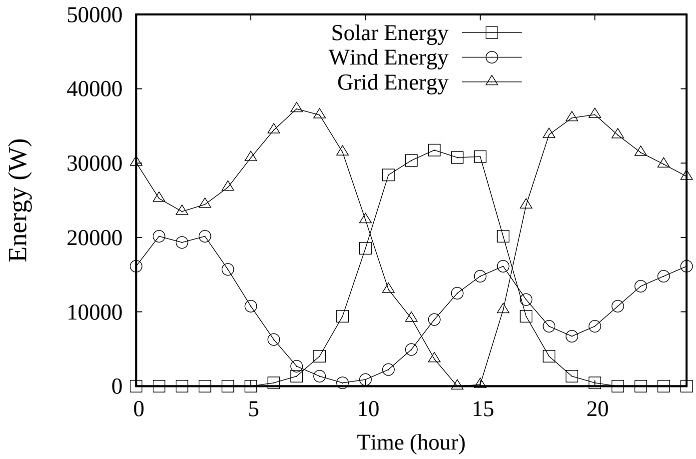

Fig. 15. Energy consumed on the data centers, differentiated for energy source.

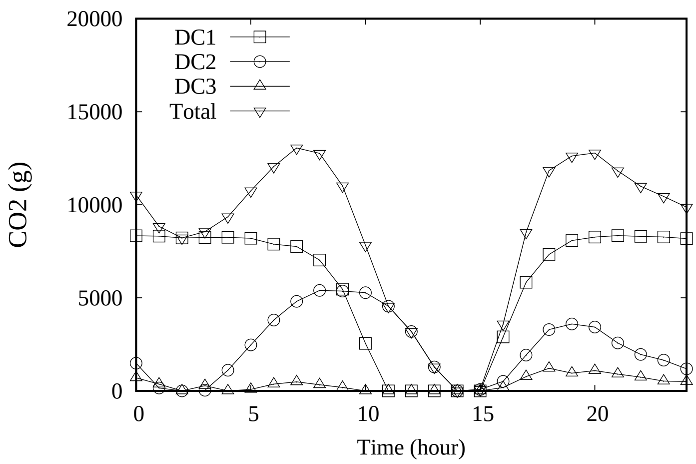

Fig. 16. Carbon emissions for the three data centers, and overall value.

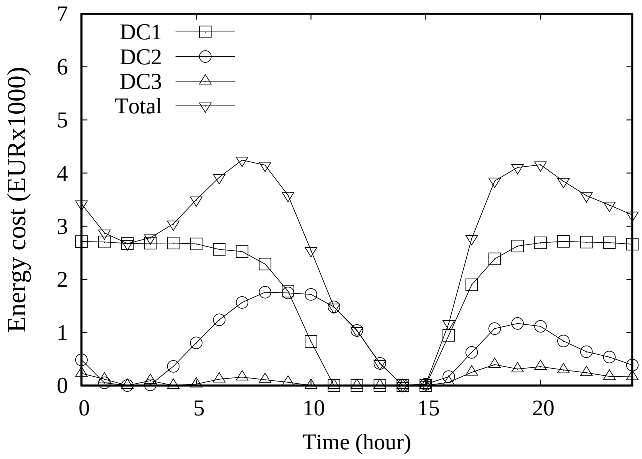

Fig. 17. Energy cost for the three data centers, and overall value.

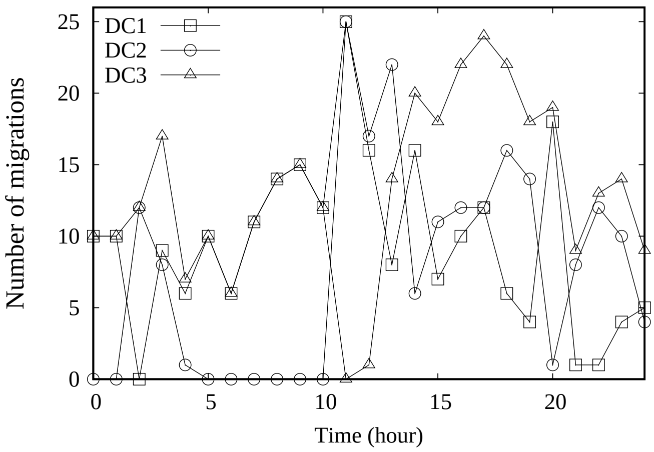

Fig. 18. Number of migrations on the three data centers.

Results with Setting 3

Setting 3 corresponds to the values α=0.4, β=0.4 and γ=0.2 in Expression (6). This setting is intermediate with respect to Setting and Setting 2. Therefore, the most relevant objectives are the reduction of carbon emissions and energy cost, but the load balance is also an objective, even if less important.

We notice that the results follow the choice of the parameters, i.e., they are intermediate between those obtained with Setting 1 and Setting 2.

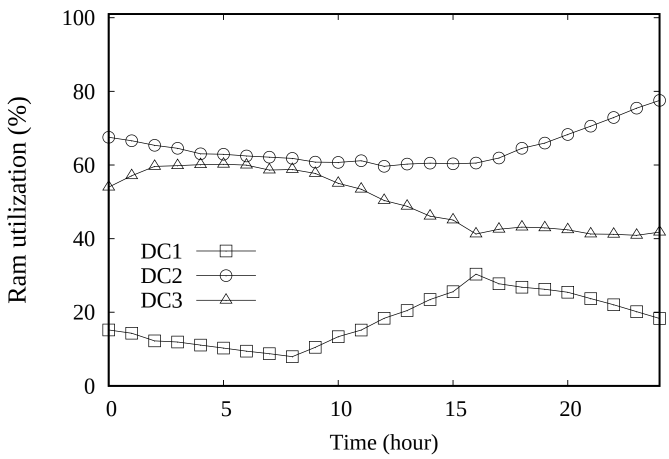

Fig. 19. RAM utilization on the three data centers.

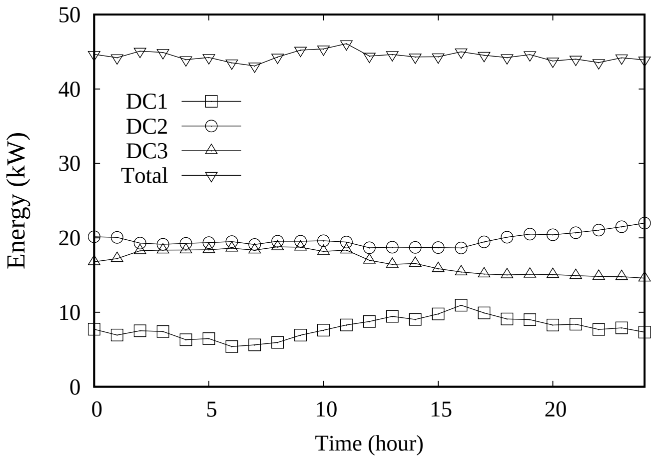

Fig. 20. Consumed energy for the three data centers, and overall value.

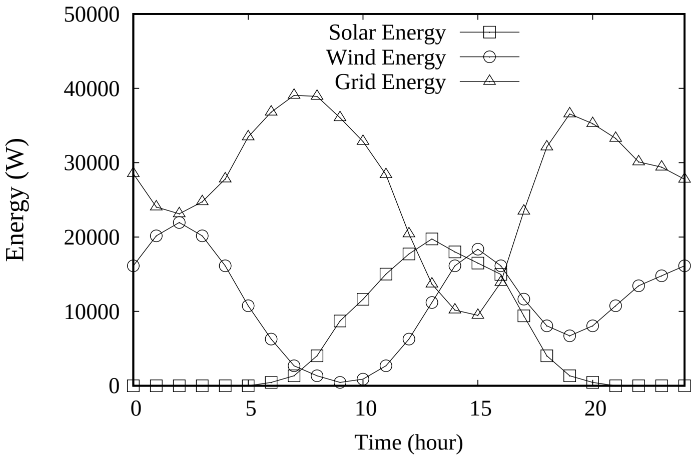

Fig. 21. Energy consumed on the data centers, differentiated for energy source.

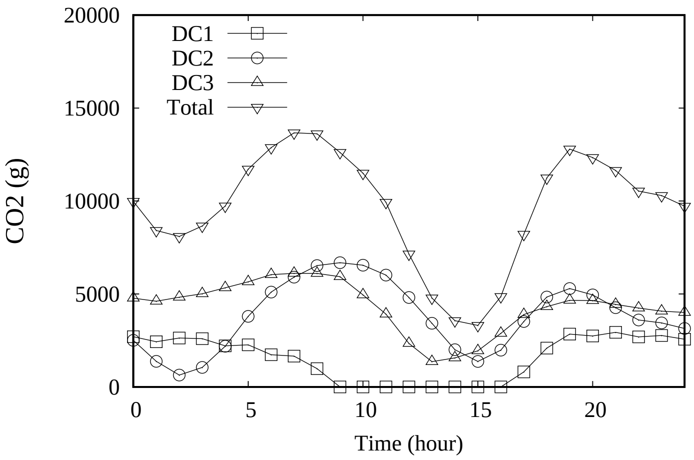

Fig. 22. Carbon emissions for the three data centers, and overall value.

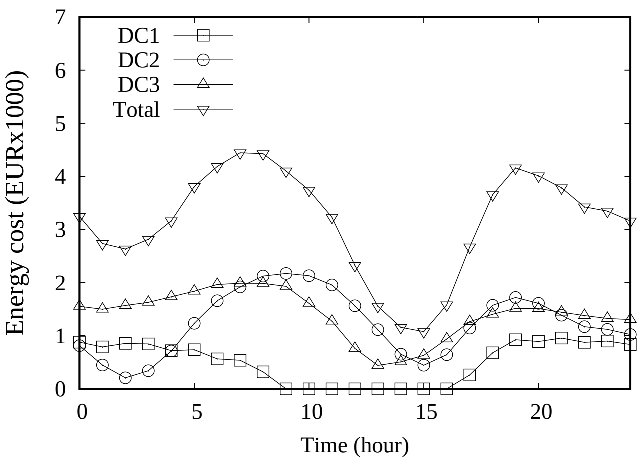

Fig. 23. Energy cost for the three data centers, and overall value.

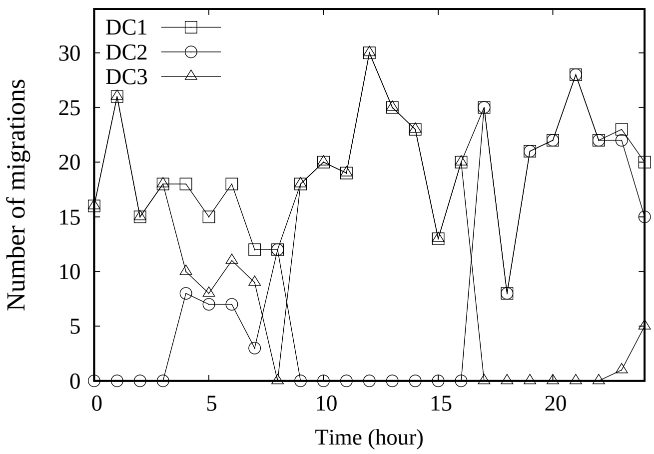

Fig. 24. Number of migrations on the three data centers.

Comparison of results with and without migrations

In this section we consider the intermediate setting of parameters, i.e., Setting 3, with α=0.4, β=0.4 and γ=0.2, and compare the results obtained with and without enabling the inter-data center VM migrations, which are the main feature of the EcoMultiCloud algorithm. More specifically, we compare the three following scenarios:

– **No migrations.** The migrations among remote data centers are disabled. However, the initial assignment of VMs is performed in accordance with the EcoMultiCloud assignment algorithm.
– **Random policy.** The migrations among remote data centers are enabled, and the choice of the VMs to migrate is made randomly.
– **Energy saving policy.** The migrations among remote data centers are enabled, and the choice of the VMs to migrate is made in accordance with an "Energy saving policy", which aims to select the CPU-intensive VMs for migration, since the energy consumption is strictly correlated to the CPU utilization.

Figure 25 shows the value of overall consumed energy per hour. We can notice the great improvement, about a 20% reduction, when the remote migrations are enabled. A small improvement, about a 2% reduction, is observed when using the energy saving policy for the selection of VMs to migrate.

Figures 26 and 27 show, respectively, the amount of green and grid energy consumption. We can notice that the reduction of energy ensured by remote migrations is mostly due to the reduction of grid energy, which of course is a very beneficial outcome. Instead, a reduction of green energy is only observed when the amount of adopted grid energy is very low, i.e., in the central hours of the day.

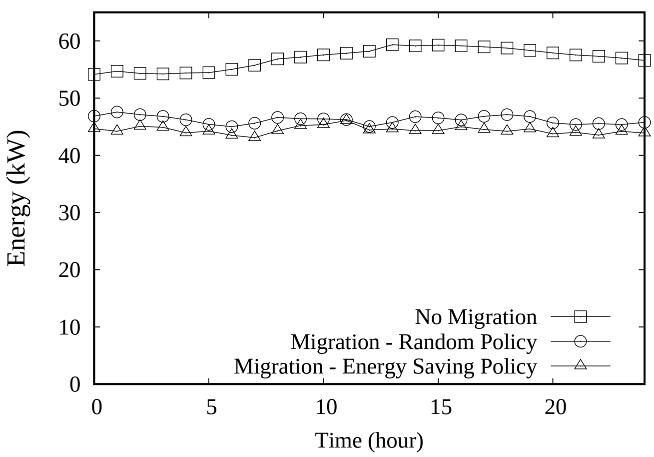

Fig. 25. Overall consumed energy with no migrations, random policy and energy saving policy.

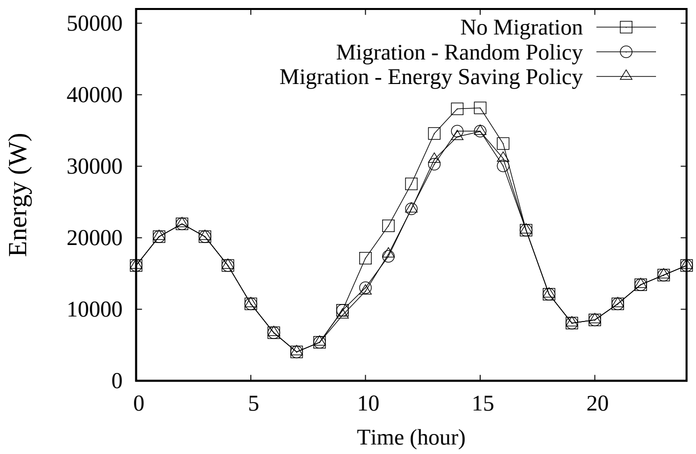

Fig. 26. Overall consumed green energy with no migrations, random policy and energy saving policy.

In Figures 27 and 28 we can see that the remote migrations allow to greatly reduce both the energy costs and the carbon emissions. The trend of these indices is strictly related to the amount of grid energy, reported in Figure 27. Indeed, we recall that the cost of green energy is zero as well as the amount of carbon emissions related to the use of green energy. In terms of energy cost, the reduction is equal to about € 31,000 per day, corresponding to more than 30%. In terms of carbon emissions, the reduction is equal to about 100 kg, corresponding to more than 30%.

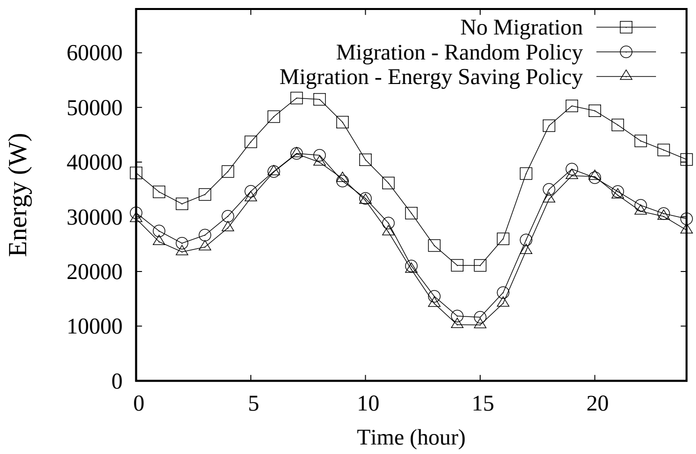

Fig. 27. Overall consumed grid energy with no migrations, random policy and energy saving policy.

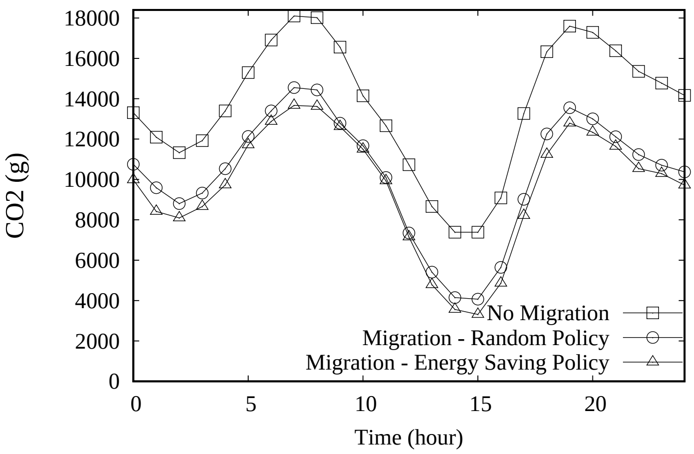

Fig. 28. Overall carbon emissions with no migrations, random policy and energy saving policy.

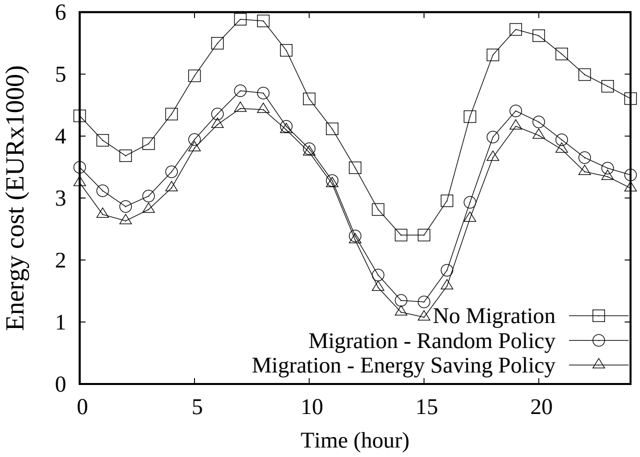

Fig. 29. Overall energy cost with no migrations, random policy and energy saving policy.

4.2 Algorithm Optimization

We presented the results of Section 4.1 to the data center administrators of TIM, former Telecom Italia. They were impressed by the ability of our algorithm of reducing the energy cost and the carbon emissions, and also by the possibility of tuning the algorithms parameters depending on the primary business and technical goals. This flexibility is considered very important, as the goals, and their relative relevance, can change with time.

We asked the DC administrators if it is possible to improve the performance, or if there is some advice from their side that can help to better match the business and technical goals. Their answer was that we should take two issues into considerations:

1. it is very important to consider the available bandwidth for remote migrations. Indeed, the inter-DC bandwidth can be used for several other purposes, for example: deploying of new applications, back up, changes in the management of different type of applications, etc.
2. it is also important to minimize the number of migrations, without affecting the other goals. This aims to improve the quality of service perceived by users, since any migration introduces some downtime and, in general, some small service degradation.

Based in the above considerations, the data center administrators and we agreed on the necessity of (i) assessing the performance of the EcoMultiCloud algorithm when varying the available bandwidth; (ii) making an additional effort to reduce the number of migrations. The first aspect was analyzed by testing several values of the inter-DC bandwidth, to inspect the impact of this parameter on the business and technical goals. The second aspect was tackled by considering an additional method for filtering the VMs to migrate. Specifically, we adopted an energy saving policy: we sorted the VMs so as to prioritize the VMs with a low ratio between disk utilization and RAM utilization, in order to minimize the amount of data to be transferred. Then we selected the VMs that use the largest amount of CPU, as this will help to move the maximum amount of workload with the minimum number of remote migrations. The results are reported in the following section.

4.3 Discussion of Results Achieved after Optimization

In this Section we present the results obtained when varying the available bandwidth for remote migrations, as advised by TIM administrators, and we tested the additional method, introduced in the previous section, for the selection of VMs. We called this method "Energy Saving Plus", and compared it to the "Energy Saving" policy adopted for the experiments commented in Section 4.1.

In Figure 30 we report the overall energy cost in the 24 hours for the entire environment. We can see that the energy cost presents small variations when varying the available bandwidth. This means that the EcoMultiCloud algorithm is effective even when the available bandwidth is reduced for any reason. We also notice that there is a small difference between the use of the "Energy Saving Plus" and the use of the "Energy Saving" policy. Indeed, the choice of the former policy aims to minimize the number of migrations and has little effect on other performance indices. Very similar considerations can be done by observing Figures 31 and 32, related to the carbon emissions and consumed energy, respectively.

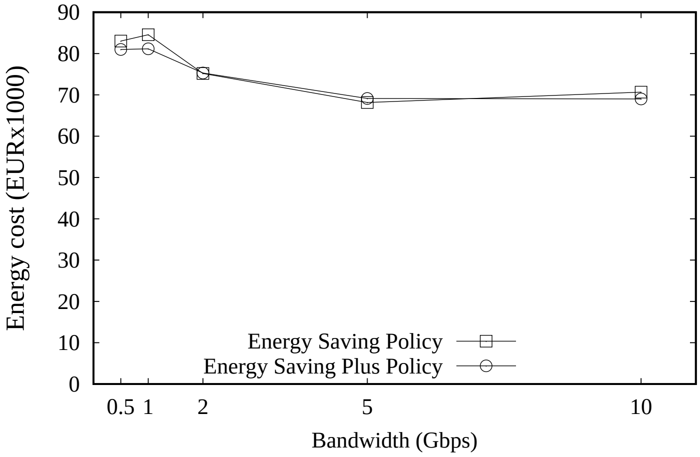

Fig. 30. 24 hours energy cost for different values of bandwidth, with and without selection of bigger VMs for remote migrations.

In Figure 33 we report the overall number of migrations performed in the 24 hours. Obviously, the number of migrations increases with the available bandwidth, as expected. When combining this result with those reported in Figures 30-32, we can conclude that an increase of the bandwidth leads to a considerable increase in the number of migrations but has a small effect on the business objectives. In other words, a relatively small amount of migrations is sufficient to achieve good performance without significant service degradations due to VM migrations. Moreover, we can notice that with the adoption of the "Energy Saving Plus" policy it is possible to achieve a significant reduction in the

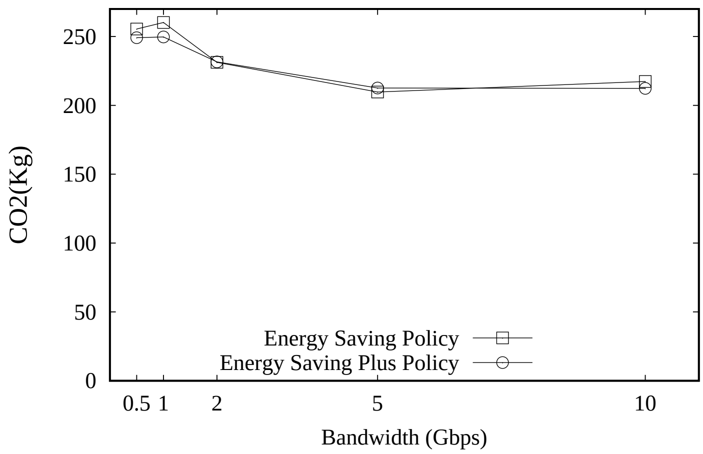

Fig. 31. 24 hours carbon emissions for different values of bandwidth, with and without selection of bigger VMs for remote migrations.

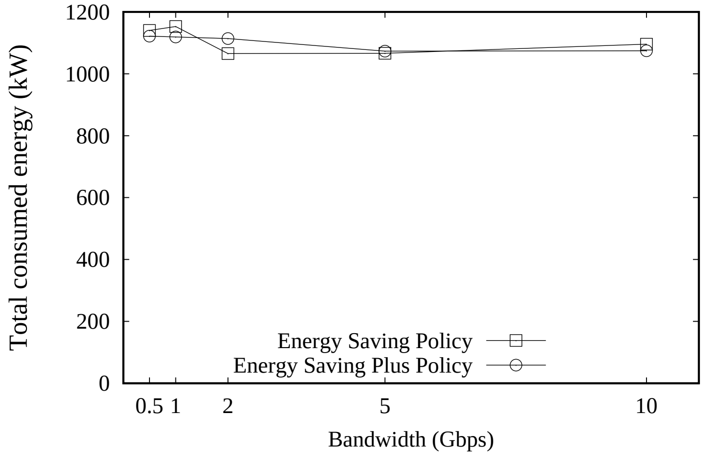

Fig. 32. 24 hours consumed energy for different values of bandwidth, with and without selection of bigger VMs for remote migrations.

number of migrations. As an example, this reduction is equal to about 30% when the available bandwidth is set to 2 Gbps.

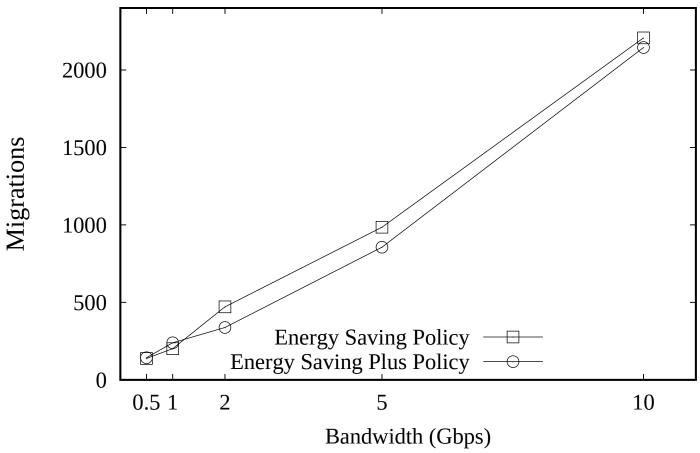

Fig. 33. 24 hours migrations for different values of bandwidth, with and without selection of bigger VMs for remote migrations.

5 Conclusions

In this document we have presented a custom solution for workload assignment and redistribution suitable for the TIM environment. The considered scenario is composed of three interconnected DCs of the TIM company across Italy. The infrastructure is distributed across Italy. For the sake of simplicity the three DCs were considered equal in architecture, number and capacity of hosts. This choice allows us to better understand the behavior of the software, specifically when the availability of green energy and the convenience of the single data centers changes with time.

The document is structured as follows: Section 2 individuates the business and technical goals that are of primary interest for the TIM company, giving details about how they are computed and evaluated. Moreover, the section describes how the algorithms for workload assignment and redistribution of the EcoMultiCloud software are customized and tuned to match the selected goals.

Section 3 describes the adopted testbed and specifies the type of energy produced in the three data centers: each data center has a peculiar green energy production based on solar or wind energy, or both. We also introduced a dynamic approach to consider the PUE of the data centers, depending on the workload and the external temperature.

Section 4 describes the obtained results. Specifically, Section 4.1 shows the values of the main performance indices, i.e., the energy cost, the carbon emissions, the amount of energy, the type of energy (green or grid), the workload distribution, and the number of migrations. The results prove that the EcoMultiCloud algorithm is very effective and also very flexible. Indeed, the administrators can decide the relative importance of the business goals by tuning the corresponding weighing parameters. Then, in Section 4.2 we summarize the feedback received by TIM administrators, aimed at optimizing the performance of the algorithm. In particular, TIM administrators suggested us to focus on the amount of available bandwidth and on the number of migrations. Following these advices, we refined the policy for the selection of the VMs to migrate, with the objective of reducing the number of migrations. In Section 4.3 we present the results achieved with the optimized algorithm, and when varying the available bandwidth for remote migrations, using the values of bandwidth indicated by TIM administrators. The results show that the performances of the algorithm are not worsened when the bandwidth is limited, which proves the effectiveness and flexibility of the EcoMultiCloud algorithm. Furthermore, we show that the optimized policy used for the selection of the VMs to be migrated, allows the overall number of migrations to be significantly reduced, up to a 30% reduction when the available bandwidth is set to 2 Gbps.

Future activities will refine this work, in order to use the data gathered in real time from a sub-set of TIM DCs. The goal is to value and tune a set of parameters that will be used in the production environment, in order to target the business and technical goals set by the TIM administrator.

References

1. Agostino Forestiero, Carlo Mastroianni, Michela Meo, Giuseppe Papuzzo, and Mehdi Sheikhalishahi. Hierarchical approach for efficient workload management in geo-distributed data centers. *IEEE Transactions on Green Communications and Networking*, 1(1):97–111, March 2017.
2. Atefeh Khosravi, Saurabh Kumar Garg, and Rajkumar Buyya. Energy and carbon-efficient placement of virtual machines in distributed cloud data centers. In Felix Wolf, Bernd Mohr, and Dieter an Mey, editors, *Euro-Par 2013 Parallel Processing: 19th International Conference, Aachen, Germany, August 26-30, 2013. Proceedings*, pages 317–328, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
3. C. Mastroianni, M. Meo, and G. Papuzzo. Probabilistic consolidation of virtual machines in self-organizing cloud data centers. *IEEE Transactions on Cloud Computing*, 1(2):215–228, July 2013.
4. Neil Rasmussen. Electrical efficiency measurement for data centers, Schneider, Whitepaper vol. 154, 2007. http://www.apc.com/salestools/nran-72754v/nran-72754v_r2_en.pdf.