An Approach to Reduce Carbon Dioxide Emissions Through Virtual Machine Migrations in a Sustainable Cloud Federation

Maurizio Giacobbe, Antonio Celesti, Maria Fazio, Massimo Villari and Antonio Puliafito
DICIEAMA, University of Messina
Contrada di Dio - 98166 Messina, Italy
Email: {mgiacobbe,acelesti,mfazio,mvillari,apuliafito}@unime.it

***Abstract*—Nowadays, Cloud federation is paving the way toward new business scenarios in which it is possible to enforce more flexible energy management strategies. Considering independent Cloud providers, each one is exclusively bound to the specific energy supplier powering its datacenter. The situation radically changes if we consider a federation of cooperating Cloud providers. In such a context, a proper migration of virtual machines among providers can lead to a global energy sustainability strategy. In this paper, we present a new strategy to reduce the carbon dioxide emissions in federated Cloud ecosystems. More specifically, we propose a solution that allows providers to determine the best green destination where virtual machines should be migrated in order to reduce the carbon dioxide emissions of the whole federated environment.**

Keywords*—*Cloud Computing, Federation, Virtual Machine Migration, Energy Sustainability, Green Computing.

I. INTRODUCTION

Cloud federation is an emerging topic in the ICT world. It refers to a mesh of Cloud providers that are interconnected through the exploitation of open standards in order to provide a universal decentralized computing environment, where everything is driven by constraints and agreements in a ubiquitous, multi-provider infrastructure [1]. From a “philosophic” point of view, the term "federation" refers to a type of system organization characterized by a joining of partially “self-governing” Clouds united by a “central government”. In a federation, each self-governing status of the involved Clouds is typically independent and may not be altered by a unilateral decision of the “central government”. Considering the Cloud computing ecosystem, besides large providers, small providers can not directly compete with mega-providers such as Amazon, Google, and Rackspace. The result is that often small/medium Cloud providers have to exploit services of mega-providers in order to develop their business logic and their Cloud-based services. Federation aims to address such an issue promoting a cooperation among small/medium Cloud providers, thus enabling on one hand the sharing of virtual resources (e.g., processing, storage, network) and on the other hand the development different hybrid scenarios [2]. This is possible by using the virtualization technology that allows providers to allocate Virtual Machines (VMs) in their physical DataCenters (DCs) and to move (or migrate) them among other federated Cloud DCs. The advantage of transforming a physical DC in a Cloud virtualization infrastructure in the perspective of Cloud federation is twofold. On one hand, small/medium Cloud providers

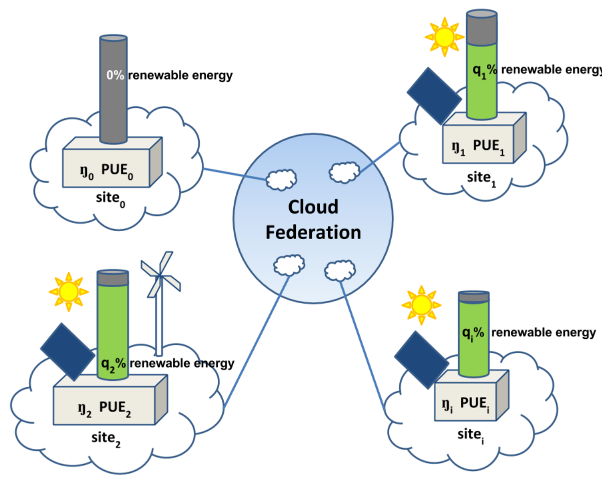

Fig. 1: Eco-sustainable federated Cloud ecosystem.

which rent resources to other providers can optimize the use of their assets, which often are under utilized, and, at the same time, earn money for the renting of their resources to other Cloud providers. On the other hand, external small/medium Cloud providers can elastically scale their logical virtualization infrastructure borrowing resources and paying them to other providers. Thanks to Cloud federation, providers can migrate virtual resources (e.g., VMs) into other providers in order to address different business scenarios, for example to deploy distributed services over a wide geographical area, to create backups in other Clouds for fault-tolerance, to move resources into secure locations, to guarantee a particular QoS, to perform software consolidation by using external Clouds in order to achieve energy cost reduction or energy sustainability.

In this paper, we focus on both Cloud federation and energy sustainability issues in order to address the reduction of carbon dioxide emissions in an eco-sustainable federated Cloud environment, as depicted in Figure 1. In particular, starting from our previous work [3], we propose a two-steps approach: the first step determines a *destination/granularity matrix* for *energy cost-evaluation* in terms of carbon dioxide emission-per-kWh. The second step elaborates the *optimum migration path* on a forecast period.

The rest of the paper is organized as follows. Section II

978-3-901882-70-8

discusses background and related works. Section III presents the main parameters that have to be considered to plan ahead energy sustainability strategies. Section IV motivates our work. Both DEST and BGD algorithms are presented in Section V. Section VI concludes the paper.

II. Background and Related Work

Energy efficiency is one of the major keywords in the Cloud computing literature [4]. In this context, the European Council has recently approved the *2030 Framework for Climate and Energy* [5]. It defines the energy and climate goals to be met by 2030, among which a binding European Union target of at least 40% reduction of greenhouse gas emissions compared to 1990, and at least 27% of renewable energy used at European Union level. The promotion of electricity produced from renewable energy sources is a high European Community priority. The purpose (art.1) of the *Directive 2001/77/EC* of the European Parliament is *"to promote an increase in the contribution of renewable energy sources to electricity production in the internal market for electricity and to create a basis for a future Community framework thereof"* [6]. Specifically, for the purposes of this Directive (art.2), renewable energy sources shall mean renewable *non-fossil* energy sources, i.e. *wind, solar, geothermal, wave, tidal, hydropower, biomass, landfill gas, sewage treatment plant gas and biogases*, therefore excluding electricity produced as a result of storage systems. Energy efficiency in Cloud computing includes two main aspects, i.e., energy costs-saving and energy sustainability that are generally uncorrelated. The first aspect aims to enforce software consolidation strategies in order to push down energy costs, whereas, the second one aims to reduce the CO_2 emissions. Although both aspects are apparently correlated the strategies to enforce them are different. In addition, finding a right compromise among them is not so trivial.

Regarding energy sustainability in Cloud federation, the Eco2Clouds Project [7] is a very important context in which groups of researchers are developing methods, guidelines and technology for enriching Cloud computing and CO_2 footprint reduction. A description of the monitoring infrastructure for eco-efficient Cloud federation and metrics are show in [8], based on the Zabbix [9] monitoring framework.
In [10] the authors present two complementary energy-efficiency optimization approaches covered in the scope of EU projects: *CoolEmAll*, with focus on building energy efficient datacenters, and *Eco2Clouds*, with focus on energy-efficient Cloud-application deployment in federated Cloud-environments. They describe metrics applied in these projects to assess and optimize energy-efficiency. In [11] the authors present a multi-objective genetic algorithm, named *MO-GA*, to optimize the energy consumption, reduce CO_2 emissions and generate profit of a geographically distributed Cloud computing infrastructure. They propose a *Pareto resource allocation* approach for Clouds focused on *energy, green house gas emission* and *profit* and use MO-GA to find the best *scheduling* according to the above mentioned goals. Their work differs from other studies in literature because it deals with both computing and energy consumption in the proposed energy model, and their approach exploits the geographical distribution in a Cloud federation.

III. Datacenter Efficiency

In this Section, we introduce the main factors related to the energy monitoring in DCs that we considered to design our solution. We considered both internal and external DC resources and the related amount of power consumption. The power consumed in a DC belonging to a federated Cloud can be expressed as follows:

$$P_{DC} = P_{IT} + P_{EE} + P_{CS} + P_{NET} \tag{1}$$

where:

- P_{IT} is the Information Technology (IT) power consumption;
- P_{EE} is the Electrical Equipment power consumption (e.g., Uninterruptible Power Supply (UPS), Power Supply Unit (PSU), Power Distribution Unit (PDU), cables, lights, batteries);
- P_{CS} is the Cooling System power consumption (e.g., Heating Ventilation and Air Conditioning and Refrigerating (HVACR) technology);
- P_{NET} is the power consumption related to the external network equipments that allow Cloud DC to move data (e.g., VMs) into other Cloud DC in a federated Cloud environment.

As motivated in [12], we assume that the energy consumption is largely independent from traffic volume and, hence, we do not consider the P_{NET} contribution in our strategy.

Following the recommendations of ASHRAE TC 9.9 and The Green Grid Association on the *Power Usage Effectiveness (PUE)* of DCs [13], i.e. the recommended metric to characterize the DC infrastructure efficiency, we express the PUE on the basis of (1), as follows:

$$PUE = \frac{P_{DC}}{P_{IT}} \tag{2}$$

Nowadays, the average PUE among the companies surveyed is lower than 1.8 [14]. Therefore, the internal energy efficiency of a Cloud DC is achieved by fully using IT resources to provide services, reducing the incidence of electrical and cooling equipment. To this end, the power consumption due to *'zombie'* IT equipment, i.e. the IT equipment that consume energy but does nothing (e.g., *'comatose'* servers), is a big problem. Moreover, we introduce a new factor η_i, that is a coefficient able to represent the energy impact on the *i-th* possible destination if we migrate a workload. It is estimated as follows:

$$\eta_i = \frac{kWh_i}{kWh_s} \tag{3}$$

For example, if we need 2 kWh_s as produced by IT energy consumptions to maintain 100 GB in the source Cloud DC, a η_i coefficient equal to 0.9 specifies that the same resources consume 1.8 kWh_i of IT energy consumption at the *ith* destination. For each federated Cloud DC, an additional carbon footprint reduction can be achieved by decreasing PUE, the introduced η coefficient or both. If a PUE reduction is difficult to be achieved for a possible destination, a reduction of the η coefficient (e.g., by using more efficient blade-centers at its DC) allows DCs to reach a significant advantage in terms of sustainability anyway.

978-3-901882-70-8 ©2015 IFIP

IV. Motivation

The objective of our approach is to enable federated Cloud DCs to enforce a dynamic energy management strategy for the whole ecosystem able to reduce the CO_2 emissions. This is possible performing a smart VM migration among federated Cloud DCs considering the best Green Destinations (GDs), i.e, powered by renewable energy sources. Our reference scenario is not static, because the energy efficiency of a DC can change in given moments of the day and in given periods of the year. In fact, the energy efficiency of renewable energy source in a given geographical area can change according to the time zone. For example, the efficiency of Photo-Voltaic (PV) system can change in different moments of the day (i.e., morning, afternoon, evening, night) and in different periods of the year (i.e., spring, summer, autumn, winter). Furthermore, due to the time zone, considering two countries in the same moment, e.g., Italy and Brazil, the energy efficiency situation is different: when in Italy is morning in Brazil is night; when in Italy the climate is cold, in Brazil it is warm. Thus if a Cloud DC is non-green in a particular period, this does not mean that it cannot be green in another period. Similar considerations can be made also for other renewable energy sources. Therefore, migrating VMs in the right "green" DCs in a dynamic eco-sustainable federated Cloud environment can reduce the CO_2 emissions of the whole ecosystem. VM migration is fundamental in our strategy. However, the most sceptic readers could ask themselves: *Does VM migration suit this reference scenario? Are there performance issues? What are the conditions to migrate VMs and where?* In the following, we briefly try to answer to all these questions and we motivate our paper. VM migration consists in transferring a VM from a source to a destination physical host. Basically, there are two kinds of VM migrations that suits different scenarios, i.e. "hot" (or live) and "cold". In "hot" migration the VM does not lose its status and the user does not perceive any change. In "cold" migrations the VM loses its status and the user notices a service interruption. The downtime is referred to the time elapsing from the instant when the VM is turned off in the source host and the moment in which the same VM is turned on in the destination host. If the downtime is negligible and the VM status is maintained, the migration is defined "hot", otherwise "cold". Considering the live migration, several works are available in the literature that aims to improve the performances.

In our opinion, Cloud federation allows providers to reduce the CO_2 emissions by means of VM migration in an energy sustainability scenario. In fact, we motivate such an idea. Federation allows Cloud providers to perform a smart software consolidation by migrating VMs from non-green Cloud DCs to green Cloud DCs in order to reduce P_{IT} and possibly part of P_{EE} and P_{CS}. A similar assumption can be made also for an energy costs-saving scenario, that is out of scope of this paper. However, this does not mean always improving the PUE value too.

V. An Approach to Reduce CO_2 Emissions in Cloud Federation

In the perspective of Cloud-to-Cloud workload migrations in an energy cost-effective federated Cloud environment, in this Section we discuss how to decide if and where moving a specific workload W (e.g., consisting of several VMs) from a Cloud DC where the energy costs is expensive to the most convenient federated Cloud DCs where the energy cost is cheaper in order to push down the energy costs of the whole federated Cloud ecosystem.

Let us consider a workload (i.e., a set of VMs belongin to the same running task) that has to be migrated from a host *a* of Cloud DC *a* into a host *b* of Cloud DC *b* in a federated ecosystem. We identify with $w_{a,b}(t)$ the portion of workload that is migrated at time *t*. Therefore, the total workload $m_{a,b}$ that is migrated from host *a* to host *b* within a time *T* is:

$$m_{a,b} = \int_{t_0}^{T} m_{a,b}(t)\, dt \tag{4}$$

However, we also need to consider a possible remaining part $r_{a,b}$ running on the source host *a* within a time *T* + τ, thus to determine $w_{a,b}$ as follows:

$$w_{a,b} = m_{a,b} + r_{a,b} = \int_{t_0}^{T} m_{a,b}(t)\, dt + \int_{T}^{\tau} r_{a,b}(t)\, dt \tag{5}$$

The remaining workload $r_{a,b}$, can be caused by possible *slowing down* in transport (e.g., due to *bottlenecks*, fall in available bandwidth) adding an extra-time τ that is necessary to conclude the migration. Thus, this extra-time can cause an additional energy consumption at source host. Therefore, we suppose the total workload W to be migrated from a source DC to a possible federated Cloud destination is:

$$W = n * w_{a,b} \tag{6}$$

where *n* is the number of source hosts. To this end, in particular, we present a two-steps solution to determine the *best green destination* in order to reduce the carbon dioxide emissions-per-kWh. Specifically, a first step is represented by the Algorithm 1 which allows to determine the *destination/granularity matrix* for *energy cost-evaluation* in terms of carbon dioxide emissions-per-kWh. Therefore, on the basis of the Algorithm 1, a second step is mainly represented by our Algorithm 2 which allows to determine the *optimum migration path* for *carbon dioxide emission reduction*. In this way, we obtain a path which is a sequence of G elements, where each one of them is the *best green destination* (identified by its index i) for the *j-th* time period. We refer to a monthly granularity G on which the forecast period is determined considering a "hot migration". In fact, the time to migrate, that is in the order of hours, is negligible if we consider a granularity of at least one month. Starting from to consider a Cloud federation characterized by N different DCs, in the first step the Algorithm 1 starts monitoring each site for the forecast period. Therefore, it evaluates the energy consumption function *cons* at each *i-th* site for a specific W that needs to be migrated. In this regard, the $cons[i]$, at line 4 of the Algorithm 1, contains the estimation of electricity consumptions in terms of kWh at the *i-th* site. It specifically depends on the IT energy consumption to maintain the workload W at the *i-th* destination, the *h* time, the energy impact η (see Equation (3)) and the PUE (see Equation (2)) as follows:

$$\phi(W, h, \eta_i, PUE_i) = W * h * \eta_i * PUE_i \tag{7}$$

where the *h* term in Equation (7) is the time (i.e., the number of hours for the days of the forecast period) spent by W in the running state (that is W is not switched as inactive).

978-3-901882-70-8 ©2015 IFIP

Algorithm 1 DEST. Destination/Granularity Matrix for Cost Evaluation

1: **while** true **do**
2:   **for** $j = 0$ to $G - 1$ **do**
3:     **for** $i = 0$ to $N - 1$ **do**
4:       $cons[i] = \phi(W, h, \eta_i, PUE_i)$
5:       **if** $available\_resource_{ij} \geq W$ **then**
6:         $dest[i][j] = \psi(cons[i], emiss_{ij})$
7:       **else**
8:         $dest[i][j] = NULL$
9:       **end if**
10:     **end for**
11:   **end for**
12:   wait(T1)
13: **end while**

Algorithm 2 BGD. Optimum Migration Path for Carbon Dioxide Emission Reduction

function path(dest)
{
  $opt[k][i] = min(dest)$
  return opt
}

For each time j, if hosting capabilities at the *i-th* site are sufficient in order to manage W (as expressed at line 5), the Algorithm 1 estimates the effective destination/granularity matrix (we name $dest$ at line 6) by the ψ function. Otherwise, it assigns a NULL value to the matrix element which is related to a destination without sufficient hosting capabilities (see line 8 of Algorithm 1). By focusing on the ψ function, it is expressed as follows:

$$\psi(cons[i], emiss_{ij}) = (\sum_{y=0}^{Y-1} (q_{ijy} * emiss_{ijy})) * cons[i] \quad (8)$$

In order to model our system, we assume that each Cloud DC uses a percentage quota of renewable energy (such as wind and solar) to power its own plant, and that it presents different features in terms of geographical location, productivity, and performance of its electrical generation system, and then of CO_2 emissions per kWh. Therefore, we consider that each site has stipulated a *Power Purchase Agreement (PPA)* contract, buying clean energy from an energy supplier. This assumption is motivated by the fact that Cloud providers can receive funds to realize new *green* plants producing clean energy and that they can receive *Renewable Energy Certificates (RECs)*. The q_{ijy} value is the percentage quota of the *y-th* electricity supply technology for the *i-th* site (e.g., the percentage quota of renewable solar and wind, rather than coal, natural gas, or nuclear) in the *j-th* time period so as determined by the set level of granularity G. The $emiss_{ijy}$ value, instead, represents the pollution emission in terms of *kg*CO_2*-per-kWh* which is related to the *y-th* electricity supply technology, for each *j-th* time period at the *i-th* site.

After determining the destination/granularity matrix for energy-cost evaluation, the Algorithm 2 is able to apply the policies to select the best GD for each time granularity. In particular, we propose to extract the destination site with minimum CO_2 emissions on the forecast period. Therefore, the Algorithm 2 stores all these values in the final $opt[k][i]$ structure thus to finally determine the *optimum migration path*. In this final structure, k is the cost of energy at the site i in terms of $kgCO_2$ and i clearly identifies the related specific destination.

VI. Conclusion

Energy sustainability represents one of the major challenges in the perspective of Cloud federation. In this paper, we presented an approach, enabling federated Cloud providers to enforce a dynamic energy management strategy for the whole ecosystem to reduce the CO_2 emissions. Specifically, the proposed approach aims to determine a *destination/granularity matrix* for *energy cost-evaluation* in terms of carbon dioxide emission-per-kWh, and the *optimum migration path* on a forecasted period.

Acknowledgment

This work was partially supported by the SIGMA Project - Italian National Operative Program (PON) 2007-2013, and by the SIMONE Project - Regional Operative Program (PO FESR) 2007-2013.

References

[1] A. Panarello, A. Celesti, M. Fazio, M. Villari, and A. Puliafito, "A requirements analysis for iaas cloud federation," in *CLOSER 2014 - Proceedings of the 4th International Conference on Cloud Computing and Services Science*, 2014, pp. 584–589.

[2] M. Fazio, A. Celesti, M. Villari, and A. Puliafito, "The need of a hybrid storage approach for iot in paas cloud federation," in *IEEE 28th International Conference on Advanced Information Networking and Applications Workshops, IEEE WAINA 2014*, 2014, pp. 779–784.

[3] A. Celesti, A. Puliafito, F. Tusa, and M. Villari, "Energy sustainability in cooperating clouds," in *CLOSER 2013 - Proceedings of the 3rd International Conference on Cloud Computing and Services Science*, 2013, pp. 83–89.

[4] L. Heilig and S. VoB, "A scientometric analysis of cloud computing literature," *IEEE Transactions on Cloud Computing*, vol. PP, no. 99, pp. 1–1, 2014.

[5] European Commission, GREEN PAPER A 2030 framework for climate and energy policies /* COM/2013/0169 final */.

[6] Directive 2001/77/EC of the European Parliament and of the Council of 27 September 2001 on the promotion of electricity produced from renewable energy sources in the internal electricity market. Official Journal L 283 , 27/10/2001 P. 0033 - 0040.

[7] The Eco2Clouds Project, http://eco2clouds.eu/.

[8] M. G. Axel Tenschert, "Cloud federation monitoring for an improved eco-efficiency," 2013.

[9] Zabbix, http://www.zabbix.com/.

[10] E. Volk, A. Tenschert, M. Gienger, A. Oleksiak, L. SisÃş, and J. Salom, "Improving energy efficiency in data centers and federated cloud environments," in *CGC*, 2013, pp. 443–450.

[11] Y. Kessaci, N. Melab, and E.-G. Talbi, "A pareto-based metaheuristic for scheduling hpc applications on a geographically distributed cloud federation," *Cluster Computing*, vol. 16, no. 3, pp. 451–468, 2013.

[12] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, "Energy consumption in wired and wireless access networks," *IEEE Communications Magazine*, vol. 49, no. 6, pp. 70–77, June 2011.

[13] R. American Society of Heating, A.-C. Engineers, and Ashrae, *PUE: A Comprehensive Examination of the Metric*, ser. ASHRAE datacom series. Amer Society of Heating, 2014.

[14] Uptime Institute® The Global Data Center Authority™, 2013 Data Center Industry Survey.

978-3-901882-70-8 ©2015 IFIP