Energy-efficiency Enablers and Operations in Software-Defined Environments

Stuart Clayman, Lefteris Mamatas[†], and Alex Galis
Dept. of Electronic & Electrical Engineering, University College London, London, UK
[†] Dept. of Applied Informatics, University of Macedonia, Thessaloniki, Greece
Emails: s.clayman@ucl.ac.uk, emamatas@uom.gr, a.galis@ucl.ac.uk

I. Introduction

New flexible network environments emerged in the last years, such as the Software-Defined Networks, the Network Function Virtualization and the various network slicing technologies. Their main goal is to satisfy and to flexibly cope with the changing requirements of both users and service / infrastructure providers with respect to the constraints of physical resources, the frequently changing network conditions, and with strategies that resolve performance bottlenecks by tuning the various performance trade-offs.

Such environments are usually being hosted in distributed Data Centers (DCs) with a large number of physical machines and the relevant services are being dynamically utilizing virtual resources, i.e. virtual machines (VMs) or virtual routers (VRs) - used interchangeably here. A very critical aspect in this context is the reduction of the energy consumption but without jeopardizing the performance of the deployed services. For example, the DCs in USA consumed about 91 billion kilowatt-hours of electricity in 2013 - double the energy required to power all the New York City households [1].

Here, we demonstrate that energy-aware manipulation of virtual resources can improve energy efficiency. Such capability should not only consider the energy consumption of the physical hosts but an estimation of the consumption at the VM or VR level as well. In our demo, we are using alternative VM placement algorithms, a number of linear and non-linear energy models and applications with variable loads in terms of CPU utilization, memory allocation and communication cost.

In the following section, we present our test-bed and demo runs. In section III we provide our conclusions.

II. Test-bed & Demo Description

In the last years, we implemented our own distributed platform for testing and evaluation of software-defined virtualized network and computing environments, called *Very Lightweight Software-Driven Network and Services Platform (VLSP)* [2], [3]. For this demo, we use *VLSP* as a software Data Center and show: (i) how energy consumption can be estimated from the physical machine to the virtual resource level; and (ii) the energy-efficient placement of VMs and the performance trade-offs. Our software has been released as open-source at [4].

In figure 1, we depict the main VLSP architecture. The platform consists of four main layers:

978-1-5090-0223-8/16/$31.00 © 2016 IEEE

- *Application Layer:* includes generic network functions or services, implemented as software components and focusing on the global behavior of the network environment. Such an example is a high-level management application defining global performance goals for the system.
- *Orchestration Layer:* acts as a global network controller and includes: (i) the *Placement Engine* which guides the optimal placement of the virtual routers; (ii) the *Monitoring Manager* providing network state data by using probes (e.g. level of congestion in a link); and (iii) the *Service Orchestrator* organizing the deployment and operation of applications over the virtual entities.
- *Distribution Layer:* is a virtual resources abstraction layer residing between the *Infrastructure* and *Orchestration Layers*. Each physical machine has a distributed *Local Controller*, receiving requests from the *Orchestration Layer* and enforcing them in the associated virtual resources of the *Infrastucture Layer*.
- *Infrastructure Layer:* the physical and virtual network resources: namely the physical servers, the virtual machines and the virtual routers. In VLSP, we created a lightweight virtual entity from the scratch with its own runtime environment and basic network protocol functionality. Our main goal is to experiment with management facilities on top of scalable virtual resources, resembling both VRs and VMs, i.e. assuming uniform management on top of integrated SDN/NFV/Cloud environments.

For this demo, we extended the *VLSP* with a number of components, such as:

- A new *placement engine*, which is the *Energy Placement Engine (EPE)*, guiding the virtual router placement onto hosts that are consuming the least amount of energy.
- A number of *energy consumption models*, both linear and non-linear, predicting the energy consumption of the physical machines and used by the *EPE*. We focused on the expressiveness of the models over diverse hardware and in dynamic environments, having as input the level of CPU, memory, and network utilization.
- A *tailor-made monitoring system* for such environments, the *Lattice* [5], collecting network and server state information used as an input from our energy models.
- *Visualization tools* depicting (i) the virtual and physical topology (figure 2), (ii) the monitored server and virtual

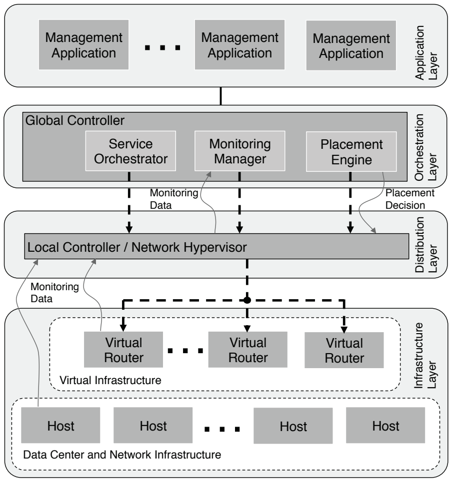

Fig. 1. VLSP architecture

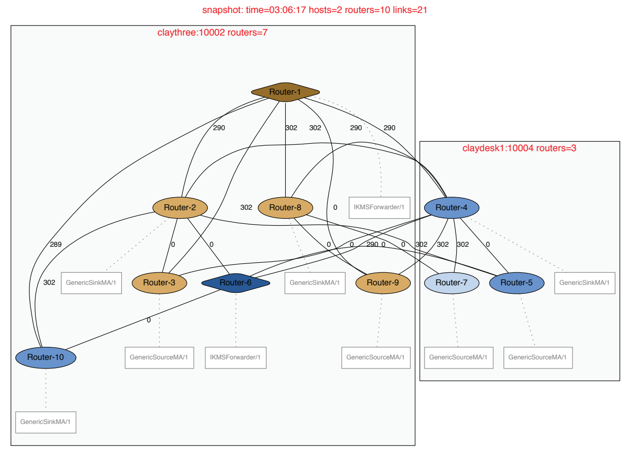

Fig. 2. Network topology visualization

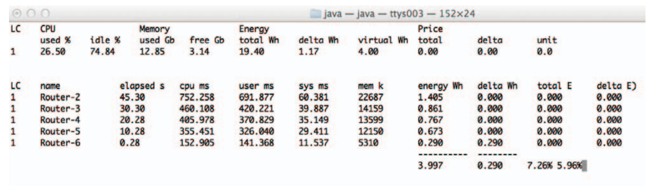

```
java — java — ttys003 — 152×24
LC   CPU                  Memory              Energy                               Price
     used %    idle %     used Gb   free Gb   total Wh   delta Wh   virtual Wh   total      delta     unit
1    26.50     74.84      12.85     3.14      19.40      1.17       4.00         0.00       0.00      0.0

LC   name           elapsed s   cpu ms    user ms   sys ms    mem k    energy Wh   delta Wh   total E    delta E)
1    Router-2       45.30       752.258   691.877   60.381    22687    1.405       0.000      0.000      0.000
1    Router-3       30.30       460.108   420.221   39.887    14159    0.861       0.000      0.000      0.000
1    Router-4       20.28       405.978   370.829   35.149    13599    0.767       0.000      0.000      0.000
1    Router-5       10.28       355.451   326.040   29.411    12150    0.673       0.000      0.000      0.000
1    Router-6       0.28        152.905   141.368   11.537    5310     0.290       0.290      0.000      0.000
                                                                       ----------  --------
                                                                       3.997       0.290     7.26% 5.96%
```

Fig. 3. Server & virtual router energy consumption & cost - UNIX `top` style

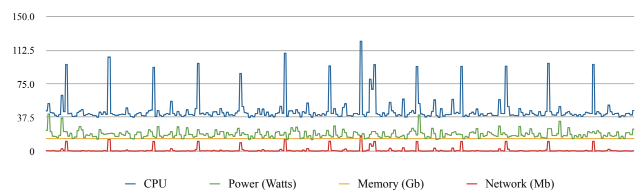

Fig. 4. Results of an energy measurement experiment

router energy consumption and cost values (figure 3) and (iii) measurement figure graphs (figure 4).

In figure 2 we show an example of two physical machines hosting 10 virtual routers. Our network visualization tool shows the virtual links with their traffic load, the deployed applications to the VRs and other useful information.

- A number of *applications with variable loads* in terms of CPU, memory and network that stress the whole facility.

Here we describe the demo run. We have one or more of physical machines hosting a number of virtual entities running distributed service nodes with different levels of CPU, memory and network load. The virtual resources are created dynamically using a probabilistic event generator and are deployed by the *Service Orchestrator* using the *Energy Placement Engine*. The EPE uses monitored values coming from the *Monitoring Manager* as input to an *Energy Model Evaluator*. The latter component estimates the energy consumption of the each entity (physical host or VM) using one of the energy consumption models. The EPE decides the placement of the virtual entities based on this estimation. The energy models are adjustable to factor the diverse hardware used. Our demo considers dynamic scenarios in terms of resource utilization, in order to demonstrate the impact of different placement algorithms, energy models and application requirements.

As the system executes, we visualize the physical and virtual network state, performance measurements, monitored values, while we show involved performance trade-offs that can be tuned. Last but not least, we estimate and visualize the energy consumption per virtual entity. This allows us to investigate the energy behavior of the system at the virtual resources level, as they are the basic units of management.

III. Conclusions

We demonstrated a novel feature of network infrastructure orchestration, the manipulation of virtual resources based on the energy characteristics of the servers and the VRs. In this context, we designed and implemented a new virtual entity placement engine, an energy model evaluator component supporting a number of linear and non-linear models, example applications with various loads and a number of visualization tools. As such, we have shown involved performance trade-offs in a DC with a *Virtual Machine (VM) Placement Engine* being aware of the energy consumption of the virtual resources.

Acknowledgements

This work is partially supported by the European Union DOLFIN [6] project of the 7th Framework Program.

References

[1] http://www.nrdc.org/energy/data-center-efficiency-assessment.asp.

[2] L. Mamatas, S. Clayman, and A. Galis, "A Service-Aware Virtualized Software-Defined Infrastructure," *Communications Magazine, IEEE*, vol. 53, no. 4, pp. 166–174, 2015.

[3] S. Clayman, L. Mamatas, and A. Galis, "Efficient Management Solutions for Software-Defined Infrastructures," *Technical Report - University College London*, 2016.

[4] S. Clayman, A. Galis, and L. Mamatas, "The Very Lightweight Software-Driven Network and Services Platform (VLSP) Open-Source Software," http://clayfour.ee.ucl.ac.uk.

[5] ——, "Monitoring Virtual Networks with Lattice," in *Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP*. IEEE, 2010, pp. 239–246.

[6] DOLFIN, "Data Centres Optimization for Energy-Efficient and EnvironmentalLy Friendly INternet (DOLFIN FP7 project)," 2013, http://www.dolfin-fp7.eu.