Data Centres Optimization for Energy-Efficient and EnvironmentalLy Friendly INternet

Funding scheme: Specific Targeted Research Projects – STREP

Co-funded by the European Commission within the Seventh Framework Programme

Project no. 609140

Strategic objective: FP7-SMARTCITIES-2013 (ICT-2013.6.2)

Start date of project: October 1st, 2013 (36 months duration)

Deliverable D5.2 DOLFIN system integration & evaluation

Due date: 30/09/2016

Submission date: 17/10/2016

Deliverable leader: IRT

Author list: Domenico Gallico (IRT), Matteo Biancani (IRT),

Dissemination Level

☒	PU:	Public
☐	PP:	Restricted to other programme participants (including the Commission Services)
☐	RE:	Restricted to a group specified by the consortium (including the Commission Services)
☐	CO:	Confidential, only for members of the consortium (including the Commission Services)

List of Contributors

Participant	Contributor
GRNET	Dimitris Siakavaras, Georgios Goumas
i2CAT	Pouria Sayyad Khodashenas, Eduard Escalona
UCL	Stuart Clayman, Alex Galis
SIEMENS	Septimiu Nechifor
PSNC	Michał Balcerkiewicz, Tomasz Ciesielczyk, Ariel Oleksiak
IRT	Domenico Gallico, Matteo Biancani
WIND	Aniello Reale, Roberto Spigolon
SYN	Artemis Voulkidis, Theodore Zahariadis
NXW	Tommaso Zini, Matteo Pardi

Amendment History

Version	Date	Partners	Description/Comments
0.1	17/05/2016	IRT	ToC definition
0.2	27/05/2016	IRT	Updated TOC
0.3	09/05/2016	SYN	Updated TOC
0.4	27/06/2016	IRT	Updated TOC
0.5	20/07/2016	SYN	Totally revised TOC
0.6	15/09/2016	SYN	Updated Evaluations with micro-DC results
0.7	16/09/2016	SYN	Updated the Evaluations (re-run the micro-DC with more iterations and added the medium-DC results)
0.8	20/09/2016	I2CAT	Inputs on section 4
0.9	22/09/2016	WIND	First (incomplete) input on section 2
0.10	23/09/2016	SYN	Input on section 3 (revenue model)
0.11	22/09/2016	GRNET	Input on section 2 regarding GRNET's DCs
0.12	27/09/2016	UCL	Scalability and Stability of the Information Flows - Evaluation
0.12.3	29/09/2016	UCL	Section 3
0.13	03/10/2016	SYN	Revised Section 3, edited Section 4 and added testbed results, edited stability of information flows results in Section 5
0.14	04/10/2016	WIND	Completed input on section 2
0.15	04/10/2016	PSNC	Section 4.3 Inputs
0.16	07/10/2016	SYN	Added Introduction
0.17	10/10/2016	SYN	Document review
0.18	13/10/2016	UCL	Reviews; added concluding remarks
0.19	13/10/2016	SYN	Minor editorial review
0.20	13/10/2016	PSNC	Section 4.3 Inputs – green power analysis
0.21	13/10/2016	WIND	Ameliorated WIND DC description
0.22	13/10/2016	SYN	Merged inputs from various partners, added lessons leasrned
0.23			
0.24	14/10/2016		Section 4.3 Inputs
0.25			
0.26	14/10/2016	UCL	REVIEWS
0.27	14/10/2016	PSNC	Chart update
0.28	17/10/2016	PSNC	Section 4.3.3 testbed diagram
0.29			
RC1			
FF	17/10/2016	IRT	Final Version

Table of Contents

Figures Summary

Tables Summary

Abbreviations

ADR	Automated Demand Response
API	Application Programming Interface
BTU	British thermal unit
CPU	Central Processing Unit
DB	Data Base
DC	Data Centre
DCO	Data Centre Operator
eCOP	energy Consumption Optimization Platform
HVAC	Heating, Ventilation, Air Conditioning
HPE	Hewlett Packard Enterprise
IaaS	Infrastructure as a Service
ICT	Information and Communications Technology
IO	Input/Output
IPMI	Intelligent Platform Management Interface
IS	Information Service
IT	Information Technology
KPI	Key Performance Indicator
LTS	Long Term Support
OpenADR	Open Automated Demand Response
OS	Operating System
PoP	Point of Presence
PUE	Power Usage Effectiveness
PV	Photovoltaics
RAM	Random Access Memory
REST	Representational state transfer
RPM	Revolutions Per Minute
SDC	Synergetic Data Centres
SGC	Smart Grid Controller

UUID	Universally Unique Identifier
UTC	Coordinated Universal Time
VCPU	Virtual Central Processing Unit
VHDD	Virtual Hard Disk Drive
VRAM	Virtual Random Access Memory
WP	Work Package
XDC	Cross-DC

Executive Summary

This document is the Deliverable D5.2 of DOLFIN project. It is entitled "DOLFIN System Integration & Evaluation" and it documents the results of integration, validation, and performance evaluations of the mechanisms implemented within the DOLFIN project.

Taking as an input all the implementation results of the components developed in WP3 and WP4 and the testbed description provided in D5.1 [1], this deliverable elaborates the validation and evaluation of the DOLFIN prototype representing the final outcome of Task 5.3. It provides a clear description of the performance and benefits that the DOLFIN prototype can offer as a solution for the efficient energy management of Data Centres.

The experiments carried out on the integrated DOLFIN platform followed the three scenarios defined in WP2. We took into account the DOLFIN DC categories as particular context for the experiments and we also have addressed the challenges identified by the project. This document presents also how the DOLFIN platform behaves in the proposed scenarios using the optimization policies implemented in the eCOP and SDC components. The analysis carried out highlights what is the estimation of the amount of energy saved and, consequently, what is the estimation on the revenue's benefits in relation to the utilization of DC resources according to the evaluation criteria described in Section 3.

In addition, the evaluation activities and results have contributed to the consortium final exploitation strategy as well as for each partner's specific exploitation plans, documented in deliverable D6.4 [2].

1 Introduction

As Information and Communication Technologies (ICT) dominate almost every aspect of our lives, the dependence of people on them also increases at ever growing rates. In this context, the use of cloud computing services has been steadily increasing, leading large Information Technology (IT) corporations to increase their relevant investments in view of increased anticipated revenues [3]. Although these giant corporations account for a great deal of the overall cloud-computing services provided worldwide, smaller-scale urban Data Centres (DCs) are responsible for almost half of the total energy consumption attributed to operating DCs [4]. Simultaneously, it is widely accepted that the perspective for achieving substantial energy savings is valid [5], simultaneously helping towards achieving environmental friendliness by means of reducing carbon emissions and maximizing the use of green energy [6].

DOLFIN proposes a coordinated yet layered approach towards achieving substantial energy efficiency gains at both individual and synergetic DCs levels. Comprising two main subsystems to achieve energy efficiency maximization at intra- and inter-DC level, namely eCOP and SDC respectively, DOLFIN has designed and built a multi-modal platform to implement the modular architecture initially presented in the deliverable D2.2 [7]. The specific design details for all the DOLFIN components are given in the various deliverables of WP3 and WP4. However, as DOLFIN should be seen a complete system able to achieve optimal energy efficiency for DCs in a synergetic, federated environment, the discrimination between the control and data flows between eCOP and SDC is no longer valid. This document presents the outcomes of the integration and evaluation efforts performed by the DOLFIN consortium in the course of delivering the integrated DOLFIN system prototype.

The outcome of the conducted integration activities is summarized in the following:

- The APIs of all DOLFIN components were harmonized against common data models, generally following the RESTful API design principles, apart from only two cases were the publish-subscribe messaging pattern was preferred to enable asynchronous operation (the SLA Renegotiation Controller and the Smart Grid Controller);
- The overall DOLFIN energy efficiency control flows were synchronized so that chained invocation of the right DOLFIN components is possible in a fully-automated manner;
- An open source VM definition and specification for automatically building a VM image containing the integrated DOLFIN components was delivered to ease deployment in different environments, ranging from cloud ones (powered by Openstack, other cloud

management platforms being also supported with minimal configuration) to simple workstations (powered by VirtualBox) in case a DC Operator would like to overview the capabilities of DOLFIN without dedicating own cloud resources. To further facilitate such system-evaluation operations by 3rd parties, we also introduced a DC-emulating component so that the capabilities of DOLFIN can be demonstrated outside production environments;

- The integrated end-to-end functionality of DOLFIN was tested against three scenarios, presented in [1], in order to showcase that DOLFIN satisfies all the requirements specified in [7] at both functional and non-functional level.

For properly evaluating the effectiveness and potential benefits of the DOLFIN system application in urban DCs, it was deemed necessary to consider the largest possible number of potential DC configurations so that the evaluation results are as subjective as possible. In this course, apart from the technical evaluation of the integration of the DOLFIN system as a whole, the DOLFIN consortium built an evaluation framework able to emulate the operation of a DC of arbitrary characteristics and assess the possible effect that DOLFIN could have in the energy consumption and the revenue of the DC, based on an SLA-aware pricing model. In this framework, we were able to test the effectiveness of the integrated DOLFIN solution on top of over 12,000 different topologies comprising IT and non-IT equipment of various computational capacity and energy consumption characteristics, also accounting for distinct average DC utilization. The results of this generic evaluation process indicated that DOLFIN could offer considerable energy efficiency results particularly in cases of relatively low average DC utilization. Notably, this case is, by far, the most common one in small- and medium-sized urban DCs as per [6].

The remainder of this report is organized as follows; in Section 2 the DC categories of interest to DOLFIN as defined by its DC-operating partners are presented, followed by a description of the evaluation criteria and methodologies of DOLFIN. Next, in Section 4, the integration activities as well as the technical evaluation of the integrated DOLFIN system in terms of end-to-end functional energy efficiency actuation are discussed. Section 5 presents the results of the generic evaluation of DOLFIN against DC topologies of variant characteristics. Last, section 6 assesses the project outcomes and presents the lessons learned from operating DOLFIN in the range of supported DC installations.

2 DC Categories in DOLFIN

In this section we will describe the four different categories of Data Centers that are represented in the DOLFIN Project. Some of the Consortium's partners (IRT, GRNET, PSNC and WIND) have made available their facilities and equipment in order to run all the tests and experiments starting from the initial phase of component's testing until the final evaluation of the integrated platform. Each of the above DC has different characteristics and provides different features; this makes all of them representatives of different categories of DC. The main differences can be associated to:

- their use (i.e., commercial and non-commercial)
- their size, number of rooms and topology
- being traditional or virtualized DC
- the utilized cooling system (i.e., free cooling, hot/cold aisle, energy reuse, renewable energy sources)

These different environments and setups contributed to improve the evaluation potential of the DOLFIN platform, in terms of flexibility and adaptability to different scenarios.

The following categories will be described in the following sub-sections:

- **Commercial distributed DC**
- **Commercial Centralized DC**
- **Non-commercial DC**
- **Non-commercial DC, with local energy generation**

2.1 Commercial distributed Data centre - Interoute

Interoute owns and operates one of Europe's largest networks and a global cloud services platform which comprises 12 data centres, and 31 colocation centres, with connections to 195 additional third-party data centres across Europe. Interoute's DCs are directly interconnected through the

Interoute's pan-European network, which spans more than 60.000 km of lit fibre as shown in Figure 2-1, below.

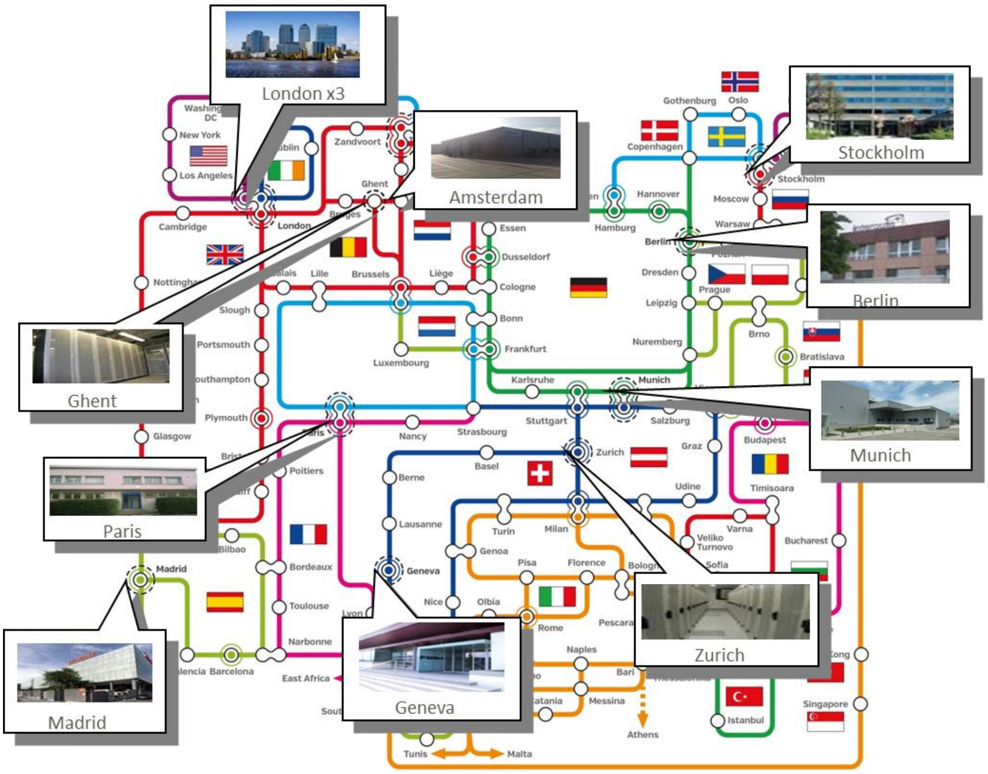

Figure 2-1: Interoute DCs

Relying on this physical infrastructure Interoute delivered Virtual Data Centre (VDC) service on January 2012, providing a scalable and customizable Infrastructure as a Service (IaaS) in a fully automated solution for on demand computing, storage and applications. The virtual infrastructure is automatically deployed and provisioned in real-time, and allows customers to specify RAM, CPU, storage, network, but also additional appliances (e.g. firewalls or IPSs) or options for scheduled local or remote backup and automated disaster recovery through the VDC Control Centre graphical interface. The automation of resource provisioning and monitoring, service recovery, intra- and inter-DC workload migration or network management is leading to a more efficient allocation of virtual instances among geographically dispersed DCs.

The DOLFIN Interoute testbed has been implemented in the co-location area of the DC facility located in Milano-Caldera as already described in D5.1. This DC, like all the Interoute's tier IV category DCs, guarantees the highest level of service availability (99.99%). This DC is fully redundant in terms of electrical circuits, cooling and network. This standard category incorporates specifications regarding the use of adequate cooling equipment as well as raised-floor system for more flexible cooling. Additionally, the standard states that cabinets and racks must be arranged in an alternating pattern in order to create "cold" and "hot" aisles. In the "cold" aisle, equipment racks are arranged face-to-face while in the "hot" aisle they are arranged back-to-back. The perforated tiles in the raised floor of the "cold" aisle allow cold air to be drawn into the face of the equipment so that this cold air washes over the equipment and is expelled out through the back into the "hot" aisle. In the hot aisle there are no perforated tiles that keeps the hot air from mingling with the cold one. The hot aisle/cold aisle cooling is showed in Figure 2-2.

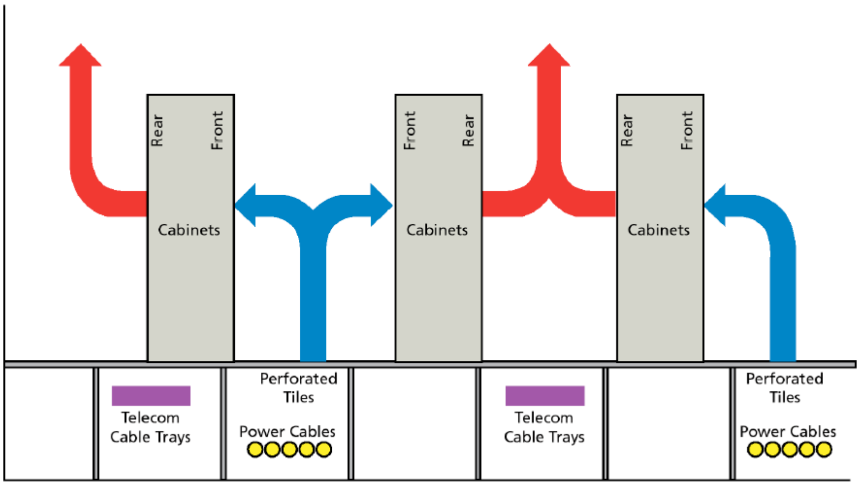

Figure 2-2: Hot aisle/cold aisle cooling

The Milan DC is a small sized DC with a gross size of 500 m^2. A portion of this space is used for a reception room and other service rooms so that the net size of customer space is approximately 300 m^2 organized in two different rooms. This customer space is completely dedicated to host approximately 240 racks with IT and networking equipment. One of the rooms of the Milan DC is showed in figure 2-3.

This facility falls in the small-medium sized DC category and as such it is can be affected by the peaks' absorption issue, so its utilization rate has to be kept quite low, in order to avoid service availability problems. Although the virtualization technology implemented allows the consolidation of workloads onto fewer servers, the overall server utilization rarely goes beyond the 40%-60%. In this context, the optimization of resources' utilization provided by the DOLFIN framework, can assume a very important role in the management of operational costs.

Figure 2-3: Interoute Milan DC, picture

2.2 Commercial Centralized Data Centre - WIND

One of the DCs considered as a target topology for DOLFIN is Wind's Ivrea DC. Although this DC form a federation with a second DC placed in Molfetta, only the Ivrea one is considered for the DOLFIN project. Therefore, from the project's point of view we can consider this as a single big centralized data centre.

Even if we are dealing with a single centralized DC, this kind of topology is important to demonstrate how the DOLFIN system may help reducing the carbon footprint of a DC also in case of no federated DCs that could help balancing the workloads. Indeed, the eCOP component of the DOLFIN system is responsible of optimizing the server utilization and therefore the energy consumption inside a single DC. WIND's DC is a good candidate in representing this topology because it is classified as Tier IV DC and is equipped with all the modern green-IT technologies, like a *free cooling* infrastructure, that make the interactions with the DOLFIN eCOP component meaningful.

2.2.1 DC Description

The DC is located in the city of Ivrea (near Turin, Piedmont Region, North of Italy) and, together with a second DC based in Molfetta (near Bari in Puglia, South of Italy), is responsible for all of Wind's business operations.

The two DCs are physically distinct and distant from each other (over 700 km) and are connected to two different Energy Providers. The facilities are designed and implemented to guarantee maximum flexibility and versatility of the DCs; the particular configuration of the federation reduces time to align infrastructure to business needs and allows Wind to:

- optimise the DC geographical distribution by HW resources utilization and manageability
- reduce incident impacts ensuring appropriate Disaster Recovery capabilities, increasing the quality of services.

The Ivrea DC follows the directives dictated by the Tiers System of the Uptime Institute. In particular, this DC can be classified as a Tier IV Data Centre: indeed, all the servers, cooling equipment, electrical and distribution facilities are redundant and dual-powered.

In the last years, the Ivrea DC has undergone a process of modernization with the goal of making it more green, thus saving on the electrical bill.

In particular, other than the ordinary modernization of the obsolete equipment, a free cooling system has been installed in the DC.

As the word *free* may suggest, the system use air coming from outside to help cooling the air circulating inside the server rooms, thus reducing the overall cooling costs. A schema representing how a general free cooling system works is depicted in Figure 2-4. We can notice how the air coming from outside enters the building in proximity of the cooling equipment used to refresh the hot air coming from the server rooms.

This configuration of the cooling system was doable thanks to the geographical location of Ivrea that guarantees cool air for the majority of the year.

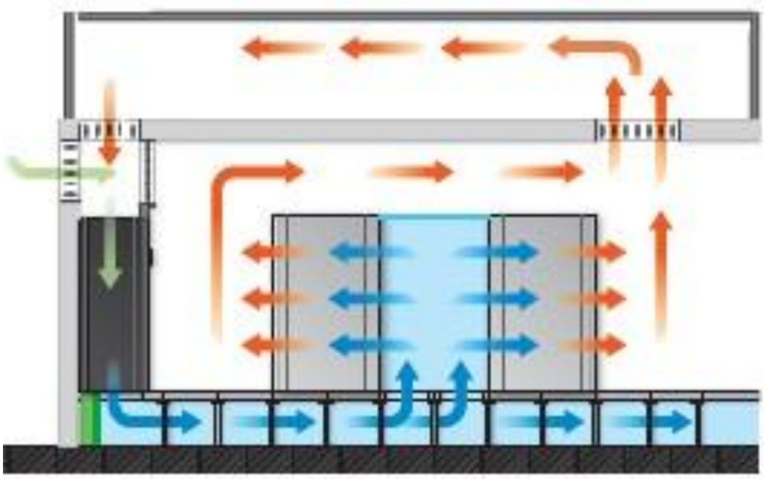

Figure 2-4: Schema of a general free cooling system

2.3 Non-commercial Data Centre – GRNET

GRNET's infrastructure consists of 3 DCs. The first one is hosted in the National Research Foundation in the centre of Athens. This DC hosts the GÉANT Point of Presence (PoP) in Athens, as well as a HellasGrid site (Grid node). It contains a total of 10 racks hosting servers and storage equipment and 14 racks with telecom equipment. Average energy consumption of the DC is 130kW. The second GRNET DC is located within the premises of the Greek Ministry of National Education and Religious Affairs in Athens. It is currently equipped with 28 racks hosting servers and storage equipment. The average energy consumption of the equipment hosted at this DC is currently around 250 kW. The third DC is a green DC that has been recently installed outdoors in the northwest part of mainland Greece, close to a power-production hydro-electric plant facility. Water from the nearby river is used to cool the equipment within the DC, while hydroelectric power for the plant facility is used for powering up the DC. The maximum power for the equipment hosted at this DC is estimated to be around 400kW and the achieved PUE is expected to be among the most competitive ones.

The DOLFIN GRNET testbed has been implemented on the second DC located at the Greek Ministry of National Education and Religious Affairs in Athens. Currently in the DC are operative 7132 logical CPUs while 1800 TB of storage space is available. This DC has been designed and implemented following high standards regarding the cooling efficiency. In-row cooling techniques are applied, with the hot/cold aisle technique (Figure 2-5). PUE is further optimized with free cooling techniques. The DC chillers are connected in parallel with air cooled heat exchangers. When the ambient air temperature drops to a set temperature, a modulating valve allows all or part of the chilled water to by-pass the existing chillers and run through the free cooling system, which uses less power and uses the lower ambient air temperature to cool the water in the system. The DC is organized in small rooms each of which resemble to small commercial DCs (Figure 2-6).

Figure 2-5: GRNET's DC in the Greek Ministry of National Education and Religious Affairs in Athens.

Figure 2-6: GRNET's testbed organization in small rooms.

2.4 Non-commercial Data centre with local energy generation – PSNC

Poznań Supercomputing and Networking Center (PSNC) is one of DOLFIN beneficiaries that operate its own DC. The new main facility opened in 2015 has currently capacity of 2MW and adopts modern Direct Liquid Cooling (DLC) techniques with heat re-use applied to the whole PSNC building – offices for more than 300 people (Figure 2-7). Additionally, PSNC still operates its former smaller DC, which is air-cooled (with some racks cooled with non-direct liquid cooling or backdoor cooling). This DC is considered as one that will be used to offer commercial hosting services. Finally, PSNC has also a micro DC setup as a part of its Laboratory of Energy Efficiency, which is used by DOLFIN and other projects as an experimental environment. Importantly, the micro DC can be connected to local renewable sources.

The main PSNC data centre is mostly used to execute complex scientific HPC workloads. Thus, PSNC has an access to the real computing infrastructure used by scientists to run their advanced applications. The centre's IT equipment includes diverse top class systems such as clusters of high performance servers, SMP machines, and hybrid CPU-GPU systems. The new main PSNC data centre consists of 1600 square meters planned for up to 180 racks and 2-16MW of power use (currently 2.5MW transformer is installed). The PSNC DC has 2 floors (+ floor with technical equipment) designed for networking equipment and low density servers (Floor 1), and HPC servers (Floor 2). Direct liquid cooling is planned for the HPC part of the data centre. Schemes of Floors 1 and 2 are illustrated in Figure 2-8:

Figure 2-7: The 360 degree view of the main server room of the PSNC data centre

Figure 2-8: Floors of the PSNC DC: Floor 1 (bottom) and Floor 2 (top)

The fastest supercomputer deployed in the PSNC data centre is the “Eagle” cluster whose computing power exceeds one peta flops (1372,13 TFLOPS). The cluster contains 32984 CPU cores of Intel Xeon E5-2697 and 120,6 TB of memory. Eagle is within 100 fastest supercomputers all over the world according to TOP500 ranking announced each year at the Supercomputing conference (SC 2015 and ISC 2016). The system organized into hot isle cabinet (see Figure 2-9) is cooled using direct liquid cooling technology from CoolIT company applied to cool CPUs and memory accompanied by in-row cooling for the remaining heat.

Figure 2-9: The Eagle computing cluster - 2nd fastest supercomputer in Poland and 80th position on TOP500 list

The PSNC facilities adopt modern ways of energy consumption reduction. In particular, the heat produced by servers through the direct liquid cooling techniques and the use of the LG VRF system cover all heating needs of the whole PSNC building. Furthermore, the micro DC can be entirely supplied from the photovoltaic system installed on a roof of the PSNC building.

In general, the PSNC computing facilities are equipped in technologies that allow reducing the energy consumption and/or its costs. This is particularly important for a public institution, which receives funds for significant investments and extensions of its infrastructure but must cover operational costs from its own funds. Currently, two main techniques are applied: a heat re-use for the whole building heating and the use of renewable energy from the photovoltaic system in the micro data centre of the Laboratory of Energy Efficiency which is the main PSNC testbed for DOLFIN. Other techniques such as thermal storage are also studied.

2.4.1 Green generation from PVs for a micro-DC installation

In recent years, in addition to efforts focused on improving energy efficiency of data centres, aspects related to energy availability and costs gained interest and importance. The reasons included increasing energy prices, problems with power grid stabilisation due to increased use of renewables and stronger incentives to reduce carbon footprint. As a consequence, researchers and data centre experts has started to invest efforts to apply renewable energy sources, energy storage, and/or automated demand response techniques to data centres both in research studies and in practice.

A good example of the use of renewable energy sources in a big non-commercial data centre is the advanced computing system Hikari [8] (Japanese for the word ‘light’) installed at the Texas Advanced Computing Center at the University of Texas at Austin, USA late August 2016 [9]. It is the first supercomputer in the US to use solar power and high voltage Direct Current (HVDC). The Hikari supercomputer cluster consists of 432 Hewlett Packard Enterprise (HPE) Apollo 8000 XL730f servers coupled with HPE DL380 and DL360 nodes that are interconnected with a first-of-its-kind Mellanox End-to-End EDR InfinBand at 100 gigabytes per second. Over 10,000 cores from 'Haswell' Xeon processors will deliver more than 400 teraFLOPS. Hikari is also a microgrid that supports a supercomputer. By day, solar panels that provide shade to a TACC parking lot also provide nearly all of Hikari's power, up to 208 kilowatts. At night, it switches back to conventional AC power from the utility grid. Another example of the use of RES in to power a DC is a study performed by MIT. Researchers at MIT has built a micro facility (a container) of 18.5 square metres being fed by a solar-PV array of 288 square metres and backed up by a utility connection and energy storage in batteries or flywheels. The use it to study how to power a data centre with renewable energy sources.

At PSNC the installation has a peak power 20kWp and consists of 80 panels (120m^2) located on the roof of the building. Additionally, energy can be stored in batteries of 75kWh capacity and, to smaller extent, using fuel cells (with maximum power supply more than 1kW). A prototype that saves energy consumed by servers using consolidation, switching off nodes, maximization of the use of renewable energy has been developed. PSNC plans to adopt the solutions created within DOLFIN to all resources of the PSNC laboratory. Details of energy and cost savings along with analysis are presented in paragraph 4.3.3.7. As successful, it might be a good proof of concept as a basis for installations for whole data centres.

2.4.2 Energy Reuse for office/space heating from the DC

The important technology allowing significant energy saving is the re-use of heat produced by servers. Especially, direct liquid cooling that enables to retrieve water of relatively high temperature, e.g. 45 degrees Celsius, can be applied for this purpose. Combining these two technologies allows reducing PUE, and consequently energy consumption of the data centre in parallel with lowering bills for heating. Additionally, it may also help in lowering operational costs. This is the case for PSNC as the energy for heating comes entirely from the data centre. Furthermore, heat at the level of MW exceeds need of the PSNC building (offices and laboratories for more than 300 people) so could be re-used somewhere else, e.g. within the campus.

3 Evaluation criteria and methodology

Based on the particular requirements documented in D2.2 [7] and considering on one hand the energy-efficiency orientation of DOLFIN and the need for generating realistic solutions that fit the needs of modern DCs (both commercial and experimental ones) on the other, the focus of the technical evaluation of DOLFIN as a whole has been on the energy- and cost-efficiency of the solutions, namely the ability of **DOLFIN to render modern Urban DCs environmentally friendlier and more sustainable**.

For the evaluation of the energy savings, apart from the actual measurements taken in DC-scope, the energy models presented in D4.1 [10] have been employed. Additionally, in order to be able to evaluate the revenue-related performance of DOLFIN simultaneously allowing for the definition and support of flexible SLAs based on performance (in terms of “greenness” and computational capabilities provided to each user), a simple revenue model has been adopted, applied in the course of the optimization procedure. The revenue model includes the calculation of the possible earnings of the DC operators due to service (computational resources) provisioning in the form of VMs and the cost that occurs due to the energy consumption of the various DC elements.

In this framework, the total revenue model that has been adopted for the evaluation of DOLFIN, assuming that the DC features S servers hosting V VMs in total and supported by a number of N non-IT infrastructure elements (e.g. lighting, HVAC), is summarized by:

$$Revenue = \sum_{v=1}^{V} \left(v.cpus \cdot v.server_{cpu\_frequency} * cpu_{mult} + v.ram \cdot mem_{mult} \right) \cdot price_{offset} - \sum_{s=1}^{S} energy\_consumption(s) - \sum_{n=1}^{N} energy\_consumption(n) \; \#(1)$$

where:

- cpu_{mult} and mem_{mult} are multipliers characterizing the contribution of the CPU and RAM usage to the price determination;
- $price_{offset}$ is a variable used for scaling the price to the current operational environment of the DC and also allowing for special pricing for individuals or special groups of users.

For the energy consumption of the servers, namely the second summation apparent in (1), the aforementioned energy models have been employed.

In order to calculate the energy consumption of the non-IT elements (lighting and HVAC), we have used generic rules and assumptions;

- It has been assumed that each rack is lit by a single lighting element of average power dissipation equal to 50W;
- It has been assumed that the energy needed to cool a server equals the BTUs of heat output it presents, using a variable server efficiency parameter ranging from 0.6 up to 0.9 [11]. We have also assumed a standard temperature difference between the outside world and the DC-internal one, so that any changes in the cooling energy consumption occurs as a function of the heat load produced by the physical servers.

Evidently, a different revenue model will result in different figures as to the performance of DOLFIN in terms of creating actual revenue out of energy efficiency or performance policy actuation (see section 5 for details and discussion). However, the chosen model has been chosen as an indicative case that enables smart and flexible SLA provisioning, based on the actual computational and energy efficiency characteristics that are being provided by the DC operators to their clients.

3.1 Evaluation methodology of DOLFIN components

The followings are the evaluation methods of the DOLFIN components:

1. Contribution & Enablers for improvements in operational efficiencies for DC operators

The following is a synthesis of the main DC energy-related operations, which are targeted by the DOLFIN components for higher efficiencies:

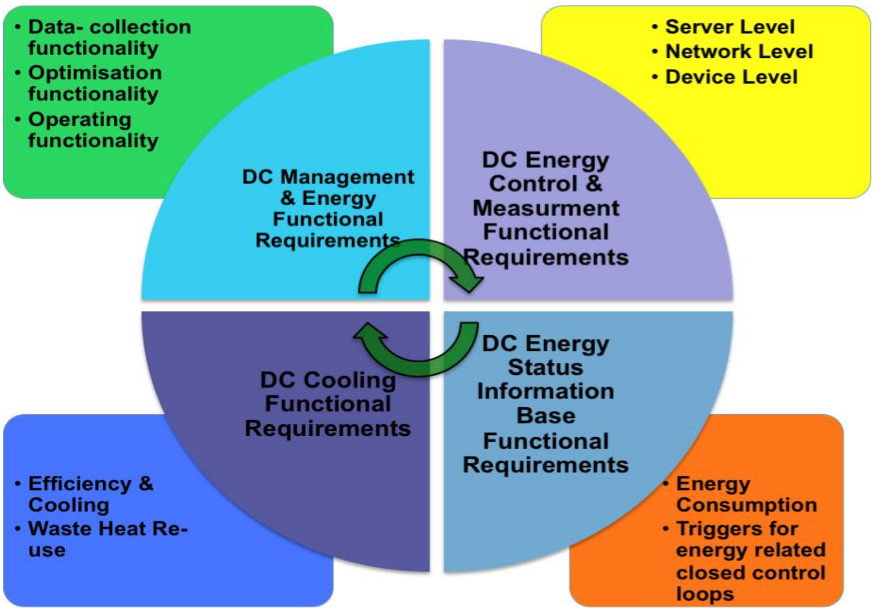

Figure 3-1: DC Operational Efficiencies

2. Contribution and Enablers for energy related DC Functionality and Characteristics

The future DCs are envisaged with the following key new characteristics targeted for the DOLFIN components:

- Continuous monitoring, energy benchmarking, dynamic control and adaptive optimisation of the Data Centre infrastructure,
- Optimisation of DC energy consumption, by dynamically changing the percentage of active/standby servers and load per server.
- SLAs, Orchestration, and Energy control loops per DC and/or group of DCs
- Optimisation of the cumulative energy consumption in a group of DCs (policy-based VMs allocation)
- Dynamic, service-effective and energy-efficient allocation of resources across a distributed network of co-operating DCs,
- Optimisation of the energy consumption at the DS / Smart City level and provide energy stabilization, by distributing VMs across a group of DCs, following the electricity demand-response approach
- Smart grid energy stabilisation, by dynamically changing the energy consumption/production requirements of DCs.
- Interconnection with the smart grid network, providing responses on the changing demands for energy, with a module that controls the legacy energy providers.

3.2 Evaluation methodology of DOLFIN as a whole

The followings are the evaluation methods of the DOLFIN system as a whole:

1. Enablement of Energy Control Loop Functionality

Future DCs would enable and employ closed control loop functionality as follows:

a) *Energy Control & Measurement Functions*, which perform Energy Control through control actions to reduce energy consumption, as specified by the *Energy Management Function*, and perform Energy Measurement by collecting measured status information. These are subdivided into Device-level, Server-level, and Network-level technologies. These functions apply to all 3 of the aforementioned hierarchy levels.

b) *Energy Management Functions*, which collect basic information, calculates the optimum case of operation, and issues operation commands to the En*ergy Control Function and* the Energy Measurement Function. This includes three sub-functions: a Data Collecting sub-function, an Optimisation sub-function, and an Operating sub-function. These functions also apply to all 3 of the DC-level, Group of Energy-conscious Synergetic DCs-level, and Smart City-level systems.

c) *Energy Status Information Base*, which encapsulates a database that stores basic information of the current mode from the Energy Control & Measurement Functions. It contains a set of status information such as energy consumption and traffic. These functions also apply to all 3 of the aforementioned hierarchy levels.

d) *Cooling Functions*, which control the efficiency of the cooling in a DC and manage the waste heat reuse.

2. Migration from separate energy control loops to a coordinated arrangement of multiple DC energy control loops

Current and future data centres are comprised of diverse cloud management and autonomic functions. The envisaged solutions accommodate the energy management with the view of:

- Improve capital and operational efficiencies for DC operators through the use of a common organization, automation, and operations of all energy functions across the different domains
- A migration from an ecosystem of separate energy management functions towards a coordinated arrangement of energy management functions as represented in the following figure.

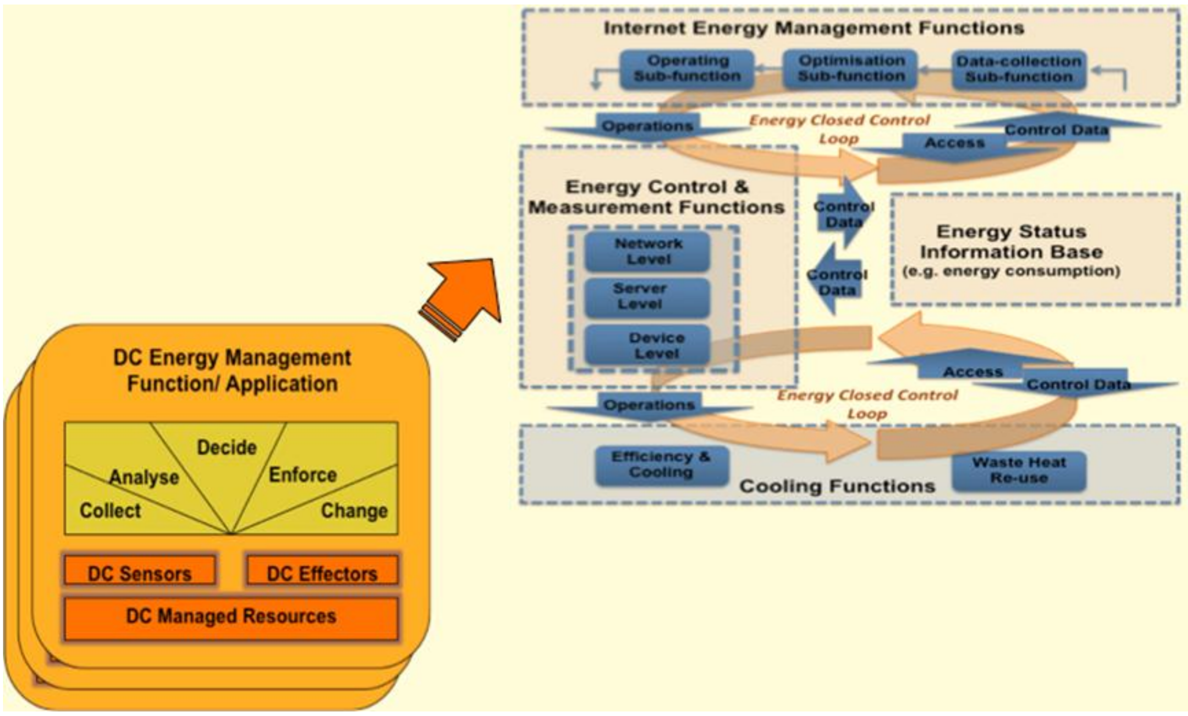

Figure 3-2: Migration from separate energy control loops to a coordinated arrangement of multiple DC energy control loops

Energy Control Loops are of prime importance in future data centres and “going green” is not only a matter of cost effectiveness and competitiveness, but also a matter of attitude and quality of life. This attitude has to permeate to all involved stakeholders, to IT system designers and developers and to managers making strategic decisions.

The main actors in green DCs who can be targeted by DOLFIN can be classified in:

- Industry
- Public organizations and regulators

- The scientific community at large (academia and research centres)
- Funding agencies
- Media contacts

With regards to industry, the group primarily includes DC operators, network and cloud operators (in most cases, these are also the service providers), and system vendors.

3. **Optimisation enablers for overall energy consumption**

Current and future data centres will monitor, and measure energy consumption and enable seamless, autonomic migration of VMs between servers of the same DC or across a group of Energy-conscious, Synergetic DCs, aiming to:

i) optimize the overall energy consumption in DCs by dynamically changing the percentage of active versus stand-by servers and the load per active server in a DC, and

ii) stabilize the Smart Grid energy distribution, under peak load and increased demand, by dynamically changing the energy consumption/production requirements of the local DCs.

iii) optimizing the energy consumption at the smart city level based on distribution of VMs across the servers that are part of a group of DCs, following an electricity demand-response approach. Enablers to feedback power to the electricity network will be employed, either by utilising electricity produced by in-house (renewable) sources or by restoring electricity from charged backup batteries.

4 Integration and testing of DOLFIN components

Offering a complete top-down solution to DC Operations management towards energy efficiency, DOLFIN employs state of the art monitoring infrastructures and advanced handling procedures of the relevant collected data as presented in the deliverables of WP3 and WP4. As data acquisition from the monitoring infrastructure is vastly different from data handling processes, in DOLFIN we separate the integration of the existing monitoring components of the trials from the integration of the various DOLFIN-enabling software components.

4.1 Integration of DOLFIN components

Since DOLFIN is modular by design [7] and most data and control flows should be synchronous by means of chained invocations of several different components, it was decided (see [12] and [10]) that all DOLFIN components should expose RESTful APIs in order to i) allow the decentralization of the DOLFIN prototype installation, ii) avoid technology lock ins and iii) facilitate integration[1]. Next, we identified four steps to allow for the seamless integration of the various DOLFIN software components:

1. Definition of API specifications;
2. Determination of common Data Models;
3. Identification of control and data flows;
4. Integration of the prototype in an automated, continuous-integration based manner.

The API reference specifications of the various components were initially designed and documented in the WP3 and WP4 implementation deliverables. In the same deliverables, the data models were also determined, primarily based on the two core DOLFIN persistence layers, namely the DOLFIN eCOP Monitor Database and the DOLFIN Information Database for the eCOP and the SDC, respectively. Similarly, the control and data flows were presented in the design deliverables [12] and [10].

[1] The SLA Renegotiation Controller and the Smart Grid Controller are both exceptions to the catholic DOLFIN directive, following the publish-subscribe messaging pattern, to support asynchronous, real-time operation.

The fourth requirement, related to the automated integration of the components following the principles of continuous integration was satisfied by means of introducing *Vagrant*, an open source platform with an aim to "*Create and configure lightweight, reproducible, and portable development environments*" [13]. In short, Vagrant allows the generation of a VM by means of simply delivering a properly formulated plain text file, following the principles of **continuous delivery** and **infrastructure automation**, one of the cornerstones of **Agile DevOps** [14]. The software configuration of the generated VM can be tuned by means of simple scripts undertaking the downloading of software dependencies, installation and configuration of specific software packages and components and their coordinated provisioning as a service. Being ideal for the purposes of a multi-cloud environment also evangelized by DOLFIN per se, the Vagrant framework is able to generate VMs that can operate on top of a variety of hypervisors including Openstack, VMWare, VirtualBox and Hyper-V as well as Docker [15], [16].

The DOLFIN Vagrant VM (already configured to support two providers, namely Openstack and VirtualBox) can be built easily by performing two simple steps:

1. Download the VM code from the relevant DOLFIN source code management platform [17];

2. Execute from command line the command

   ```
   $ vagrant up
   ```

   for building a VirtualBox-based VM, or

   ```
   $ vagrant up --provider=openstack
   ```

   for building an OpenStack-compatible VM image and launching it in a local Openstack deployment. More information on how to build and configuration the DOLFIN Vagrant VM may be sought in [17].

After running the above-mentioned commands, the following steps are performed automatically by the DOLFIN partners' pre-configured Vagrant environment:

1. An Ubuntu Linux 14.04 LTS image gets downloaded from the internet to constitute the OS supporting all the VM operations;

2. The latest version of the source code of the various DOLFIN modules get downloaded inside the VM;

3. The software dependencies of all DOLFIN modules get resolved, downloaded and installed inside the VM;

4. All DOLFIN modules get automatically configured and built;

5. All DOLFIN modules acquire service operation properties.

6. All DOLFIN modules start operating.

Updating to the latest version of the upstream code when a VM has already being built is also easy and can be performed by means of issuing the two following commands:

```
$ git pull --recurse-submodules
$ vagrant reload --provision
```

The first command makes sure that the latest VM and DOLFIN components code gets downloaded and the second command re-builds the VM.

The configuration of the VM characteristics (e.g. number of virtual CPUs or total amount of VM memory) can be tuned by simply editing a plain text configuration file. Similarly, **the configuration of the various DOLFIN components is also performed in single plain text file** (other than the VM-oriented one mentioned in the lines above). Notably, the core set of configurations related to the inherent DOLFIN operation is provided **pre-configured to work out-of-the-box without necessitating any user intervention**; the only configuration needed from the user standpoint is related to configuring the Openstack- and metring-infrastructure related information (e.g. the Identity Service endpoint URLs, the pair of valid credentials with access to the Openstack-related information, the IPMI characteristics of the servers etc. – see also §4.2) so that the various DCO Brokers are able to seamlessly integrate the existing infrastructure. In this perspective, **a single point of entrance for the configuration of the whole DOLFIN solution** is supported.

Apart from facilitating the DOLFIN Integrated VM provisioning and maintenance, a simple web dashboard allowing for the status overview and management of the various DOLFIN modules has been also made available by means of the open-source software Supervisor [18], allowing for starting/stopping/restarting the core DOLFIN services as well as monitor their health and output in real time, without necessitating logging into the VM or introducing the need for executing OS-specific commands.

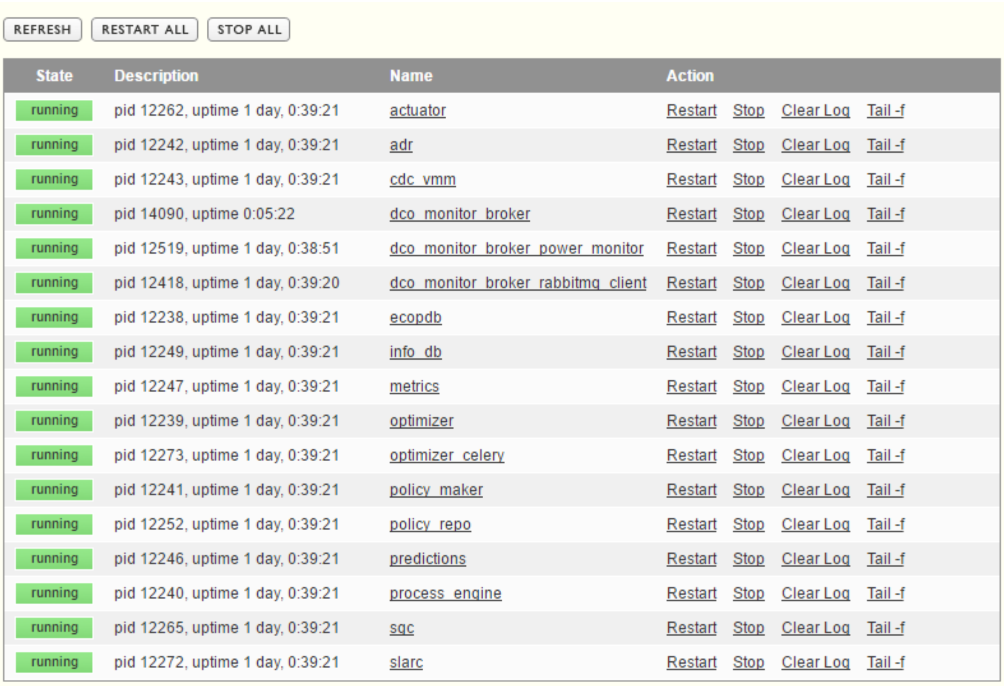

REFRESH | RESTART ALL | STOP ALL

State	Description	Name	Action
running	pid 12262, uptime 1 day, 0:39:21	actuator	Restart Stop Clear Log Tail -f
running	pid 12242, uptime 1 day, 0:39:21	adr	Restart Stop Clear Log Tail -f
running	pid 12243, uptime 1 day, 0:39:21	cdc_vmm	Restart Stop Clear Log Tail -f
running	pid 14090, uptime 0:05:22	dco_monitor_broker	Restart Stop Clear Log Tail -f
running	pid 12519, uptime 1 day, 0:38:51	dco_monitor_broker_power_monitor	Restart Stop Clear Log Tail -f
running	pid 12418, uptime 1 day, 0:39:20	dco_monitor_broker_rabbitmq_client	Restart Stop Clear Log Tail -f
running	pid 12238, uptime 1 day, 0:39:21	ecopdb	Restart Stop Clear Log Tail -f
running	pid 12249, uptime 1 day, 0:39:21	info_db	Restart Stop Clear Log Tail -f
running	pid 12247, uptime 1 day, 0:39:21	metrics	Restart Stop Clear Log Tail -f
running	pid 12239, uptime 1 day, 0:39:21	optimizer	Restart Stop Clear Log Tail -f
running	pid 12273, uptime 1 day, 0:39:21	optimizer_celery	Restart Stop Clear Log Tail -f
running	pid 12241, uptime 1 day, 0:39:21	policy_maker	Restart Stop Clear Log Tail -f
running	pid 12252, uptime 1 day, 0:39:21	policy_repo	Restart Stop Clear Log Tail -f
running	pid 12246, uptime 1 day, 0:39:21	predictions	Restart Stop Clear Log Tail -f
running	pid 12240, uptime 1 day, 0:39:21	process_engine	Restart Stop Clear Log Tail -f
running	pid 12265, uptime 1 day, 0:39:21	sgc	Restart Stop Clear Log Tail -f
running	pid 12272, uptime 1 day, 0:39:21	slarc	Restart Stop Clear Log Tail -f

Figure 4-1: Overview of the DOLFIN services administration panel.

The automation of the maintenance procedures streamlined by the use of the DOLFIN VM infrastructure significantly facilitated all integration activities as changes could be performed in any VM running from the Integration infrastructures of DOLFIN to the local development environments of the various DOLFIN partners, limiting bug discovery times and contributing to fast issues resolution through real-time, safe, “sandboxed” testing procedures.

4.2 Integration of existing infrastructures

Since continuous monitoring is the cornerstone of modern DCs operation based on the merits of extensive infrastructure virtualization, the DOLFIN solution has to be able to cope with a variety of underlying infrastructures, including monitoring equipment and cloud management platforms. DOLFIN addresses these requirements by employing the various DCO Brokers.

In the DOLFIN context, all trial testbeds (with the exception of VLSP which is a virtualized, software-defined DC and is, also, natively supported by DOLFIN DCO Brokers) are relying on Openstack to abstract and manage their physical IT resources (see deliverable D5.1 [1] for details and discussion related to the DOLFIN trial testbeds). In this framework, the DCO Brokers contained in the DOLFIN VM presented in the previous paragraph offer complete integration for all the relevant operations of interest including VMs management actions such as pausing, unpausing, scaling, migration and relocation.

Regarding physical resources monitoring and control, to support all trials, the DCO Brokers offer out-of-the-box support for:

- IT equipment supporting IPMI (Intelligent Platform Management Interface, [19]) for retrieving information related to power consumption and status, system fan speeds, temperature etc. Notably, IPMI is supported by more than 200 computer system vendors including Intel, Cisco, Dell, Hewlett-Packard, NEC etc. [20].
- Simple Network Management Protocol (SNMP)-enabled resources for retrieving information about power consumption, status etc. of IT and non-IT equipment;
- Openstack Telemetry service (Ceilometer) metrics for accessing information related to the status and characteristics of the VMs and, when available, the energy consumption of the servers in case the Kwapi [21] is supported.

The following table summarizes the set of data available in each of the three demonstration sites along with the means of acquiring this data.

	WIND Testbed	
Measurement type	Monitored entity	Monitoring Mechanism
Ceilometer Metrics	VMs	Ceilometer
Fan Speed (RPM)	Servers	SNMP
Metrics	DC	Calculated
Power (W)	Servers	SNMP
Power (W)	Racks	Calculated
Temperature (C)	Servers	SNMP

Voltage (V)	Servers	SNMP
Interoute Testbed		
Measurement type	**Monitored entity**	**Monitoring Mechanism**
Ceilometer Metrics	VMs	Ceilometer
Metrics	DC	Calculated
PSNC Testbed		
Measurement type	**Monitored entity**	**Monitoring Mechanism**
Ceilometer Metrics	VMs	Ceilometer
Fan Speed (RPM)	Servers	IPMI
Metrics	DC	Calculated
Power (W)	Servers	IPMI
Power (W)	Racks	Calculated
Temperature (C)	Servers	IPMI
Voltage (V)	Servers	IPMI

Note that in all trial sites, all Openstack Telemetry measurements for VMs are available and collected, as presented in [22].

4.3 Integration testing of DOLFIN components

As formulated in D5.1 [1], testing scenarios are hypothetical situations with an emphasis on the assessment of DOLFIN component(s). Normally, a test scenario features five key characteristics: a) a complete story that is b) motivating, c) credible, d) complex and e) easy to evaluate [23]. At each testing scenario, DOLFIN platform is used to preserve optimality of the energy consumption, by following a non-probabilistic series of actions and producing a verifiable outcome, indicating that the DOLFIN subsystems all interworked as they should and the energy consumption is minimized at aggregate DOLFIN level (depending on the active policies of the DOLFIN DCs supported).

In this section three main integration testing scenarios are presented:

- Intra DC optimization testing scenario, which evaluates DOLFIN capabilities in the context of a single DC. In this scenario the DC will be considered as a "solo eco-system" able to react to internal changes to reach the optimal energy consumption.
- SLA testing scenario, which introduces the capabilities related to DOLFIN inter-DCs cooperation in a federated DC's group, to share resources and reach the optimal energy consumption, while preserving the contractual SLA with the customers. The objective is to move the VMs between DCs while respecting specific SLA constraints.
- Smart Grid testing scenario, which is used to evaluate the DOLFIN DC adaptation capability when integrated in a Smart Grid environment. In this case the DC energy optimization logic could be directly affected by the information provided by the Utilities through the Smart

Grid interfaces. The testing scenario is used to evaluate how DOLFIN reacts over changes to energy costs or energy availability.

4.3.1 Intra DC optimization testing scenario

This scenario tests the basic self-adaptation and optimization capabilities of DOLFIN in the context of a single DC operation. Although DOLFIN is designed to support networks of synergetic DCs operating in concert, a single DC scenario is a good show case to emphasis on two main objectives pursued in the framework of the project, i.e. the optimization of energy consumption and ensuring DC user experience by defining and re-negotiating SLAs. More specific, in one hand, the testing scenario highlights the DOLFIN capabilities on reducing the DC energy consumption, which leads to less environmental footprint as well as reduced total cost of operation/use for the DC operators/consumers. On the other hand, with the help of dynamic SLA negotiation, the testing scenario demonstrates how DOLFIN platform can minimalize the performance degradation during the energy consumption optimizations and captures the effect of price incentives offered by DCs as a result of reduced anticipated power consumption. In brief, the testing scenario will provide a good understanding on the following aspects of the DOLFIN platform:

- Monitoring and managing IT and non-IT infrastructures
- Accessing user information and providing accounting/billing services
- Monitoring the power efficiency of the DC, as calculated by means of sets of well-defined measurements, and
- Optimizing the operation of the DC in terms of energy consumed

To demonstrate the above-mentioned objectives and functionalities of the DOLFIN components involved, the testing scenario takes into account both user-initiated and DOLFIN-initiated actions. It is worth noting that, the user-initiated actions are used to facilitate the testing scenario and include the insertion of artificial load to the DC infrastructure. DOLFIN actions refer to the asynchronous actions initiated by the various DOLFIN components. The basic scenario description is as following:

The DOLFIN optimization policy is set to minimize the energy consumption of the DC in absolute terms. At a certain time, the load of a particular set of VMs running on different servers is rapidly and unexpectedly increased (but can be accommodated by the DC itself). After one hour, the load is reduced to normal levels. The DOLFIN platform should identify the load changes and reconfigure the DC load allocation to the servers/racks/rooms so that in both cases its energy consumption is as minimal as possible.

4.3.1.1 Requirements addressed by the testing scenario

This testing scenario addresses the DC energy state optimization. From a technical perspective, it correlates to any situation that needs DC energy state optimization. The presented setup in this section is based on UC 1.1 stated in D2.2 [7]. In summary, the testing scenario highlights that when the overall DC efficiency reaches a given threshold a re-organization is triggered to boot the efficiency again. The expected outcome of this showcase is to reduce operational costs, by reducing the total energy consumed, while respecting the SLAs.

The DOLFIN components under evaluation in this scenario are primarily the core components composing eCOP, namely:

- ICT Performance & Energy Supervisor
- Energy efficiency policy maker and actuator
- eCOP Monitor DB

In addition to the eCOP core components, the various DCO Brokers interfacing the underlying infrastructures are subject to be evaluated as well, so as to demonstrate the ability of DOLFIN to integrate various virtualization and cloud management systems:

- DCO Hypervisor Manager
- DCO Monitor/Collector
- DCO Appliance Manager
- Monitoring Backend

Moreover, the SLA Renegotiation Controller is required to interface with the Energy efficiency policy maker and actuator group of components.

4.3.1.2 Test prerequisites

To successfully handle the test scenario the following prerequisites are needed:

1. A valid Openstack [24] installation, managing the DC resources.
2. A DOLFIN instantiation comprising all relevant components identified as test components.
3. Proper monitoring equipment should have been deployed to monitor the performance and characteristics of the DC elements of interest, including physical servers, server racks, DC rooms, HVAC equipment and lighting.

4.3.1.3 Testing setup and configuration

The following setup configuration should be performed to guarantee the performance of the testing scenario:

- A vanilla OpenStack installation is considered, so a default OpenStack configuration is assumed. Following the most minimalistic approach, we assume that at least the following services are configured: Identity (Keystone), Compute (Nova), Network (Neutron), Image (Glance) and Telemetry (Ceilometer).
- The rest of the eCOP components are configured according to the instructions provided in D3.3 [25] and D3.4 [26].

- The DCO Brokers are configured to properly mediate the OpenStack installation, HW equipment and the eCOP DB components.

It was also assumed that the Openstack defaults regarding RAM and CPU over commissioning have been kept intact, assuming a value of 1 for the memory and 16 for the CPU.

The scenario has been configured with the following starting conditions:

Table 4-1: DC Rooms considered in test 4.3.1.

ID	DC	FLOOR	NAME
1	DC-1	1	Room-1

Table 4-2: DC Racks considered in test 4.3.1.

ID	NAME	ROOM
1	Rack	1

Table 4-3: DC Physical hosts considered in test 4.3.1.

Serial Number (Openstack Hypervisor ID)			
Name	node-2.cefrielstack.com	node-3.cefrielstack.com	node-4.cefrielstack.com
Rack_id	1	1	1
Active	True	True	True
CPU	16	16	8
Ram (MB)	64308	64308	32052
Hdd	7917	7917	7917
Cpu_frequency	2400	2400	2000
Min_watt_per_cpu_core	4.63	3.75	4.75
Max_watt_per_cpu_core	10.00	11.00	30.00
Min_watt_per_kbps	0.00	0.00	0.00
Max_watt_per_kbps	0.00	0.00	0.00
Cpu_energy_mult	0.80	0.80	0.80
Ram_energy_mult	0.20	0.20	0.20
Net_energy_mult	0.00	0.00	0.00
Green	False	False	False

Table 4-4: VM Flavours considered in test 4.3.1.

ID	NAME	VCPUS	VRAM	VHDD
1	m1.tiny	1	512	1
2	m1.small	1	2048	20
48747012-ce1f-4d4c-9c22-13afe0640524	demo.medium	2	2048	4
695a4bbc-b814-4349-ad06-4c19cc9e0b4b	test	1	1024	0
a9cfc8db-9975-42a3-a180-d2654b6952b0	demo.large	2	2048	10
b6e70307-d96f-4961-b4ac-50228148abd0	demo.small	1	1024	4
e056dc69-92ab-410d-8230-0d87c6c5896d	m1.medium	2	4096	40
ecafa6e9-ae9b-44b0-a6d2-3faebc72d760	m1.micro	1	64	0

Table 4-5: VMs considered in test 4.3.1.

UUID	NAME	USER	FLAVOUR	SERVER	STATUS
21d2a378-eb21-4426-959e-3a67aeee26a3	SLARC_test	31b103874a434a6d84c448dfd6ed538b	48747012-ce1f-4d4c-9c22-13afe0640524	00301d396f5765c2dd52e1b42ebb6d94d472180d1450e751d3cb02dc	ACTIVE
2b3e39b7-be57-4e57-aea7-a63b599307fb	VM-100	8ecdac2c3b75420d879ea12b640c2656	2	67803f3dac187e481efe06671bce5532a7dd15a840526cda1a490683	ACTIVE
516c2284-79bd-4862-b8ab-a9d721e24b39	test	99b5f6ac67f34f10a00d36e0ed2b9e6e	ecafa6e9-ae9b-44b0-a6d2-3faebc72d760	3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17	ACTIVE
394bd2e5-3ff8-4fe2-8c41-79395576e255	VM-100	8ecdac2c3b75420d879ea12b640c2656	2	3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17	ACTIVE
a684a835-eb8b-4f70-9afe-bfc3aa3a2563	VM-100	8ecdac2c3b75420d879ea12b640c2656	2	3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17	ACTIVE
ab899d09-8f5c-4354-	DOLFIN	a0915f6b8f63411aa3096dc	2	00301d396f5765c2dd52e1	ACTIVE

8ccf-739ba290e5ce		2e3b229d1		b42ebb6d94d472180d1450e751d3cb02dc	
dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c	dolfin	a0915f6b8f63411aa3096dc2e3b229d1	2	00301d396f5765c2dd52e1b42ebb6d94d472180d1450e751d3cb02dc	ACTIVE

4.3.1.4 Test execution and expected results

The steps to be performed in order to perform a successful test execution are presented in D5.1 [1] and are, hence, omitted in the present document for reasons of brevity.

Granted the above configuration, we would expect all VMs to be consolidated into the most efficient physical node, that is server named node-2.cefrielstack.com, as under high load it features the least energy consumption (lowest value for the Max_watt_per_cpu_core parameter); granted the RAM and CPU over commissioning values of the Openstack environment, the node node-2.cefrielstack.com has enough resources to host all VMs even at 100% CPU and RAM usage. As some servers were already hosted in this node, the following VMs should be consolidated:

- 21d2a378-eb21-4426-959e-3a67aeee26a3
- 2b3e39b7-be57-4e57-aea7-a63b599307fb
- ab899d09-8f5c-4354-8ccf-739ba290e5ce
- dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c

In this framework, the step-wise checkpoints to acknowledge the operation of the test are as follows:

1. The Policy Maker has changed the default policy to 'local_only' and also informed the Optimizer about it;
2. The Policy Maker requested an optimization plan from the Optimizer;
3. The Optimizer devised an optimization plan containing:
   a. Live Migration directives were issued by the Optimizer regarding the VMs that do not reside in node-2.cefrielstack.com;
   b. Server hibernation directives were issued by the Optimizer regarding the physical nodes node-3.cefrielstack.com and node-4.cefrielstack.com.
4. The Optimizer requested approval from the Policy Maker in order to send the devised optimization plan to the Policy Actuator;
5. Upon approval, the Optimizer sent the Plans to the Policy Actuator;

6. The Policy Actuator performed the actions requested and logged the actions into the eCOP Monitor DB.

4.3.1.5 Outcome of the test

The following screenshots present an overview of the outcome of the test. In short, all VMs have been consolidated in the physical node with name node-2.cefrielstack.com. Also, all relevant actions have been correctly logged into the eCOP Monitor DB.

Virtual Machines

Search

UUID	Name	User	Flavour	Server	Status
21d2a378-eb21-4426-959e-3a67aeee26a3	SLARC_test	31b103874a434a6d84c448dfd6ed538b	2 CPUs, 2048 MB RAM, 4GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE
2b3e39b7-be57-4e57-aea7-a63b599307fb	VM-100	8ecdac2c3b75420d879ea12b640c2656	1 CPUs, 2048 MB RAM, 20GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE
394bd2e5-3ff8-4fe2-8c41-79395576e255	VM-100	8ecdac2c3b75420d879ea12b640c2656	1 CPUs, 2048 MB RAM, 20GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE
516c2284-79bd-4862-b8ab-a9d721e24b39	test	99b5f6ac67f34f10a00d36e0ed2b9e6e	1 CPUs, 64 MB RAM, 0GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE
a684a835-eb8b-4f70-9afe-bfc3aa3a2563	VM-100	8ecdac2c3b75420d879ea12b640c2656	1 CPUs, 2048 MB RAM, 20GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE
ab899d09-8f5c-4354-8ccf-739ba290e5ce	DOLFIN	a0915f6b8f63411aa3096dc2e3b229d1	1 CPUs, 2048 MB RAM, 20GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE
dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c	dolfin	a0915f6b8f63411aa3096dc2e3b229d1	1 CPUs, 2048 MB RAM, 20GB HDD	node-2.cefrielstack.com (3223259eb0dfb07e258857dd721560e5b4717306178016ed14360c17)	ACTIVE

Showing 1 to 7 of 7 rows

Figure 4-2: The VMs status as perceived by the eCOP Monitor Database after the execution of test 4.3.1..

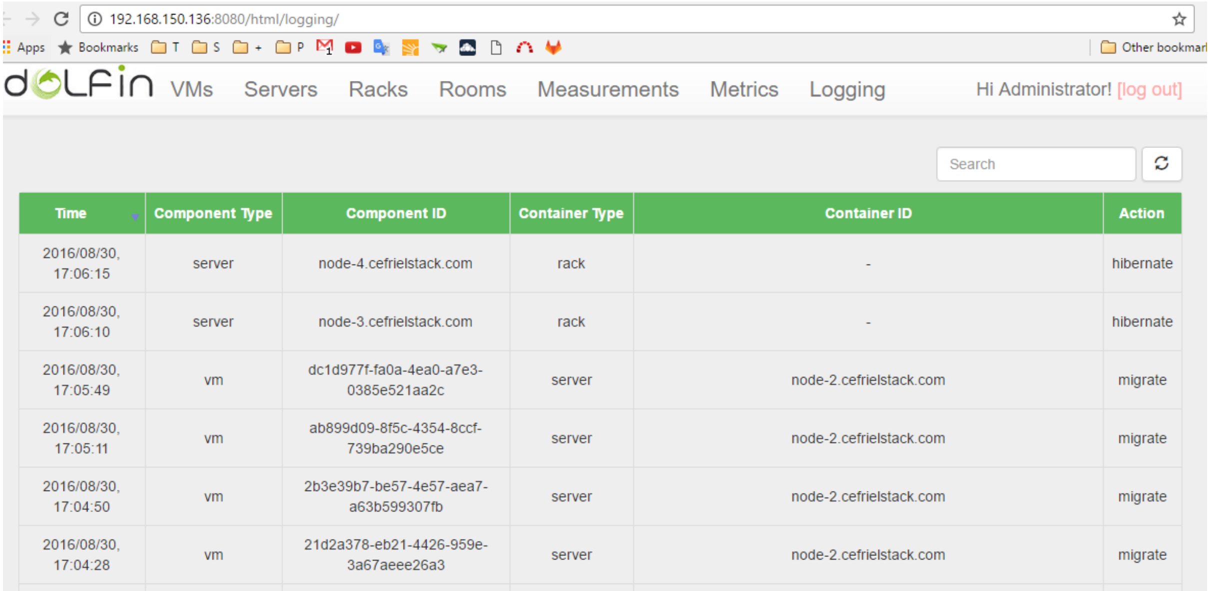

Time	Component Type	Component ID	Container Type	Container ID	Action
2016/08/30, 17:06:15	server	node-4.cefrielstack.com	rack	-	hibernate
2016/08/30, 17:06:10	server	node-3.cefrielstack.com	rack	-	hibernate
2016/08/30, 17:05:49	vm	dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c	server	node-2.cefrielstack.com	migrate
2016/08/30, 17:05:11	vm	ab899d09-8f5c-4354-8ccf-739ba290e5ce	server	node-2.cefrielstack.com	migrate
2016/08/30, 17:04:50	vm	2b3e39b7-be57-4e57-aea7-a63b599307fb	server	node-2.cefrielstack.com	migrate
2016/08/30, 17:04:28	vm	21d2a378-eb21-4426-959e-3a67aeee26a3	server	node-2.cefrielstack.com	migrate

Figure 4-3: The actions communicated by the Policy Actuator to the eCOP Monitor DB after their execution.

4.3.1.6 Testing scenario check points

The Policy Actuator performed the actions requested and logged the actions into the eCOP Monitor DB.

4.3.1.6.1 The Policy Maker has changed the default policy to 'local_only' and also informed the Optimizer about it.

After the execution of the test, the active policy of the Optimizer was as depicted in Figure 4-4. As can be overviewed, the active policy had "Energy" as target, pushing the optimizer to optimize against energy efficiency. Moreover, the set of constraints only included that no IT load should be accepted from (hence also relocated to) another (synergetic) DC.

Considering the above, this checkpoint was verified.

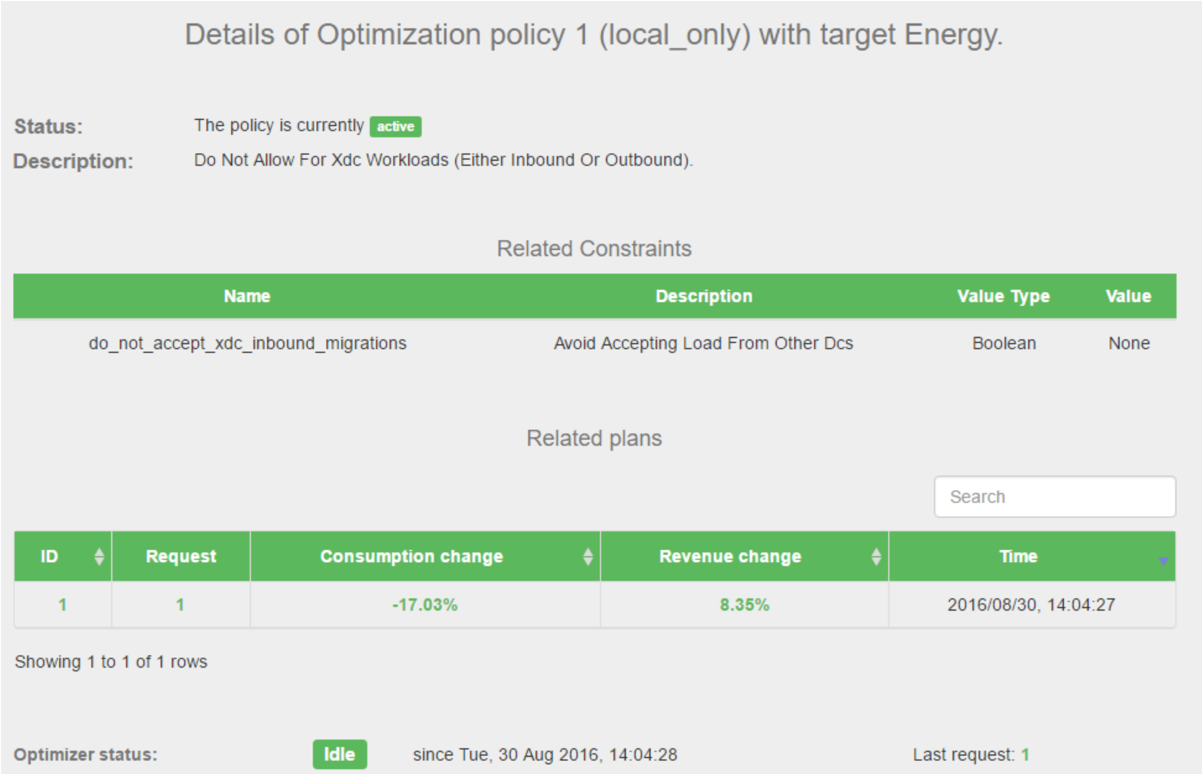

Figure 4-4: The active policy of the DC as acknowledged by the optimizer during the test 4.3.1..

4.3.1.6.2 The Policy Maker requested an optimization plan from the Optimizer

After the Optimization process ended, the view of the past Optimization plans was as follows:

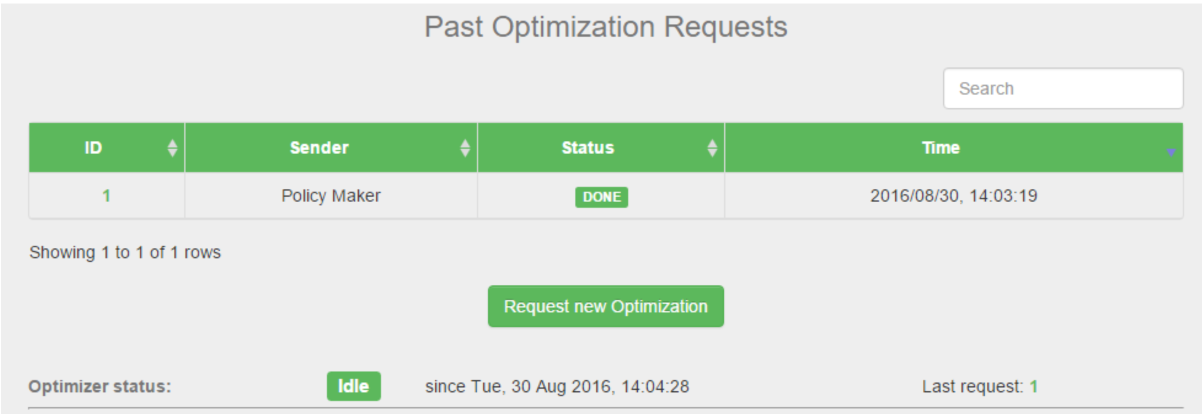

Figure 4-5: The request performed by the Policy Maker, in the context of test 4.3.1..

As can be easily deduced, the Policy Maker successfully issued the optimization command to the Optimizer.

Considering the above, this checkpoint was verified.

4.3.1.6.3 The Optimizer devised an optimization plan consolidating all VMs in one server, instructing the hibernation of the rest.

After the optimization process, the view of the relevant view of the generated plan #1 was as follows:

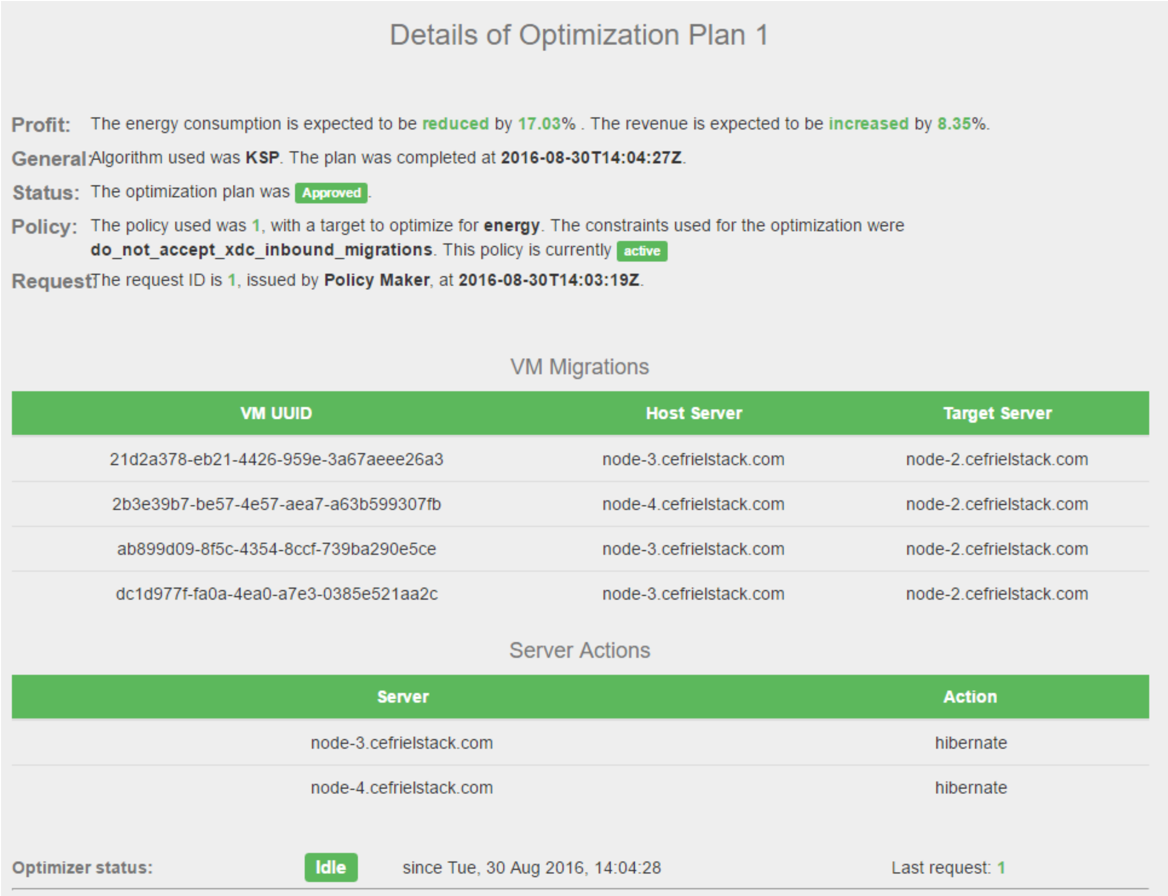

Figure 4-6: Details of Optimization Plan #1, in the context of test 4.3.1..

As can be deduced by Figure 4-6, the VMs that were expected to migrate to node-2.cefrielstack.com were issued relevant optimization directives. Also, the (after the execution of the plan empty) physical nodes were issued commands to be hibernated. It can also be seen that according to the models governing the optimization processes of the Optimizer, the expected benefit in terms of energy consumption reached 17.03%, whereas the corresponding revenue gains reached 8.35%.

Considering the above, this checkpoint was verified.

4.3.1.6.4 The Optimizer requested approval from the Policy Maker in order to send the devised optimization plan to the Policy Actuator

As can be seen in Figure 4-6, the generated plan has been characterized as "Approved". Moreover, when checking at the logs of the Policy Maker component, the following lines could be found:

```
[D 160830 14:04:28 optimizer:14] received optimizer request id: 1

[I 160830 14:04:28 web:1908] 200 POST /v1/policy/optimizer/approve (127.0.0.1)
1.58ms
```

Indicating that the Policy Maker received the optimization plan #1 from the Optimizer, then got a request for optimization plan approval.

Considering the above, this checkpoint was verified.

4.3.1.6.5 Upon approval, the Optimizer sent the Plans to the Policy Actuator;

The fact that the Policy Actuator was contacted by the Optimizer can be deduced by the fact that the actual VM migrations and the server hibernations actually took place and were logged into the eCOP Monitor DB as per Figure 4-2 and Figure 4-3. Moreover, one could find the following lines in the SLA Renegotiation Component:


```
Message received: {"payload":{"name":"node-2.cefrielstack.com-21d2a378-eb21-4426-959e-3a67aeee26a3","eventType":"STOP_VM_CRASH","id":"21d2a378-eb21-4426-959e-3a67aeee26a3"},"message":"event_notification","timestamp":"1472565890472"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-21d2a378-eb21-4426-959e-3a67aeee26a3","eventType":"START_VM","id":"21d2a378-eb21-4426-959e-3a67aeee26a3"},"message":"event_notification","timestamp":"1472565890481"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-2b3e39b7-be57-4e57-aea7-a63b599307fb","eventType":"STOP_VM_CRASH","id":"2b3e39b7-be57-4e57-aea7-a63b599307fb"},"message":"event_notification","timestamp":"1472565911740"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-2b3e39b7-be57-4e57-aea7-a63b599307fb","eventType":"START_VM","id":"2b3e39b7-be57-4e57-aea7-a63b599307fb"},"message":"event_notification","timestamp":"1472565911745"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-ab899d09-8f5c-4354-8ccf-739ba290e5ce","eventType":"STOP_VM_CRASH","id":"ab899d09-8f5c-4354-8ccf-739ba290e5ce"},"message":"event_notification","timestamp":"1472565949306"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-ab899d09-8f5c-4354-8ccf-739ba290e5ce","eventType":"START_VM","id":"ab899d09-8f5c-4354-8ccf-739ba290e5ce"},"message":"event_notification","timestamp":"1472565949310"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c","eventType":"STOP_VM_CRASH","id":"dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c"},"message":"event_notification","timestamp":"1472565969978"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-2.cefrielstack.com-dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c","eventType":"START_VM","id":"dc1d977f-fa0a-4ea0-a7e3-0385e521aa2c"},"message":"event_notification","timestamp":"1472565969983"} On Topic: vim With routingKey: key.manager and Content type: text/plain
```

```
Message received: {"payload":{"name":"node-3.cefrielstack.com-node-3.cefrielstack.com","eventType":"STOP_HOST","id":"node-3.cefrielstack.com"},"message":"event_notification","timestamp":"1472565975303"} On Topic: vim With routingKey: key.manager and Content type: text/plain

Message received: {"payload":{"name":"node-4.cefrielstack.com-node-4.cefrielstack.com","eventType":"STOP_HOST","id":"node-4.cefrielstack.com"},"message":"event_notification","timestamp":"1472565975519"} On Topic: vim With routingKey: key.manager and Content type: text/plain
```


Indicating that the component was informed by the Policy Actuator about the time that each of the migrated VMs were down during the migration process and that the physical hosts node-3.cefrielstack.com and node-4.cefrielstack.com were hibernated.

Considering the above, this checkpoint was verified.

4.3.2 SLA testing scenario

This testing scenario evaluates the ability of DOLFIN to optimize the placement of VMs within the federated DC architecture while respecting the user-defined SLAs. Since guaranteeing end-user Quality of Experience (QoE) during the attempt to minimize power consumption is an important issue for DC operators, the current testing scenario tries to highlight how DOLFIN considers user-defined SLAs to maintain the expected performance from the end-user standpoint while optimizing DC power consumption.

In this testing scenario, a set of periodic actions under normal operating conditions is simulated, to examine the number and performance of all active DC VMs (in terms of SLA compliance). Once the process of matching SLA requirements to DC availability is complete, a plan for energy optimization is produced, which is the expected outcome of this test scenario. In detail, the SLA Renegotiation Controller (SLARC) receives regular updates on the states of all existing VMs. In this way, it can assure that the SLA requirements are met. To simulate an outage or a system failure, we forcibly set a VM to enter the SUSPENDED, i.e. not ACTIVE status. After a few seconds, SLARC should notify the Policy Maker that an SLA breakage is about to happen. In turn the Policy Maker notifies the Optimizer to check about the SLA status of the particular VM and find a proper plan to fix the problem.

Note that the Green SLAs are handled considering the exactly same workflow, so the check of a single SLA type breakage is enough for guaranteeing normal operation under both possible SLA breakage scenarios.

A successful test scenario completion should be able to demonstrate:

- All involved eCOP components are able to communicate and interwork;
- All involved SDC components are able to communicate and interwork;
- The eCOP is able to coordinate with the SDC.
- DOLFIN is able to identify VM state changes and react upon them;
- DOLFIN is able to identify optimal states of DC operation as dictated by an administrator-set DC optimization policy without breaking the SLAs of the users, even if this involve SLA renegotiation;

DOLFIN is able to manage VMs to accomplish the optimal allocation plans produced by the Energy efficiency policy maker and actuator components.

4.3.2.1 Requirements addressed by the testing scenario

This testing scenario primarily illustrates DOLFIN's ability to perform SLA renegotiation. Moreover, it is a good show case for the energy efficient workload redistribution and situations with multi tariffs from the Utility companies.

The DOLFIN components under examination in this scenario are:

- ICT Performance & Energy Supervisor
- Energy efficiency policy maker and actuator
- eCOP Monitor DB
- SLA Renegotiation Controller

4.3.2.2 Test prerequisites

As this scenario also involves intra-DC optimization, all relevant elements from the Setup of the first testing scenario (subsection 4.1) are required, including the SLA Renegotiation Controller.

4.3.2.3 Testing setup and configuration

The testing setup and the configuration of the DC was exactly the same as in 4.3. The VM of interest is (see Table 4-6):

Table 4-6: The VM considered in the framework of test 4.3.2.

UUID	NAME	USER	FLAVOUR	SERVER
21d2a378-eb21-4426-959e-3a67aeee26a3	SLARC_test	31b103874a434a6d84c448dfd6ed538b	48747012-ce1f-4d4c-9c22-13afe0640524	00301d396f5765c2dd52e1b42ebb6d94d472180d1450e751d3cb02dc

Regarding SLARC configuration, the default alarm-triggering duration for the Availability-related SLA was set to 60 seconds.

4.3.2.4 Test execution and expected results

To simplify the setup and enable easy test results validation without depending on external factors, the test execution steps have been slightly changed from the initial description in D5,1 [1]. The changes pertain mostly to the initialization phase, where, in practical terms and to remove the dependence on the prediction engine results, a VM was manually set to SUSPENDED state so as to emulate a forced shut down of its operation e.g. due to a time-shifting optimization plan directive. The rest of the changes are related to the fact that the VM Priority Classifier was integrated within the Optimizer context, reducing the necessary steps to perform an optimization procedure.

Step #	Test Action	Expected Results	Means of verification
1	A command to set a VM into SUSPENDED state to emulate a forced shutdown is issued through the Policy Actuator	The VM enters the SUSPENDED state	The eCOP Monitor Database presents the VM as SUSPENDED.
2	The Policy Actuator notifies	SLARC gets notified about	The SLARC logs contain evidence

	the SLARC of the change in the VM state	the VM state change	of the SLARC notification
3	After 60 seconds, the SLARC emits an alarm that the VM SLA is about to be broken	The Policy Maker receives the alarm	The Policy Maker logs contain evidence that the alarm was caught.
4	The Policy Maker changes the DC Policy to optimize the SLA of the VM in hand and notifies the Optimizer	The Optimizer acknowledges the SLA-related policy as the active one	The Policy Maker logs contain evidence that a new policy has been issued to the Optimizer; The Optimizer UI also contains the same evidence.
5	The Policy Maker issues an optimization request to the Optimizer	The Optimizer receives an optimization request from the Policy Maker	The Policy Maker logs contain evidence that a new optimization request has been issued to the Optimizer; The Optimizer UI also contains the same evidence.
6	The Optimizer schedules a new optimization to wake the VM in hand up	The generated Optimization plan contains an action to wake the VM up	The Optimizer UI offers a plan with the updated Policy, indicating that the VM should be waken up
7	The Optimizer requests approval from the Policy Maker to forward the plan to the Policy Actuator	The Policy Maker approves the plan	The Optimizer logs the plan as "Approved", whereas the logs of the Policy Maker indicate that this transaction actually took place.
8	The Optimizer forwards the plan to the Policy Actuator	The Policy Actuator wakes the VM up	Policy Actuator logs indicate receipt of the plan.
9	The Policy Actuator wakes the VM up	The VM changes state to "ACTIVE"	The eCOP Monitor DB displays the VM as "ACTIVE" and the wakeup action is properly logged.

Table 4-7: Test steps for testing scenario 2.

As anticipated, the expected result is that the VM gets awaken soon after the relevant SLARC alarm.

4.3.2.5 Outcome of the test

Based on the setup outlined in 4.3.1.3 (Table 4-5) and considering the test steps presented in 0, the following screenshots provide evidence of the initial, intermediate and final VMs state.

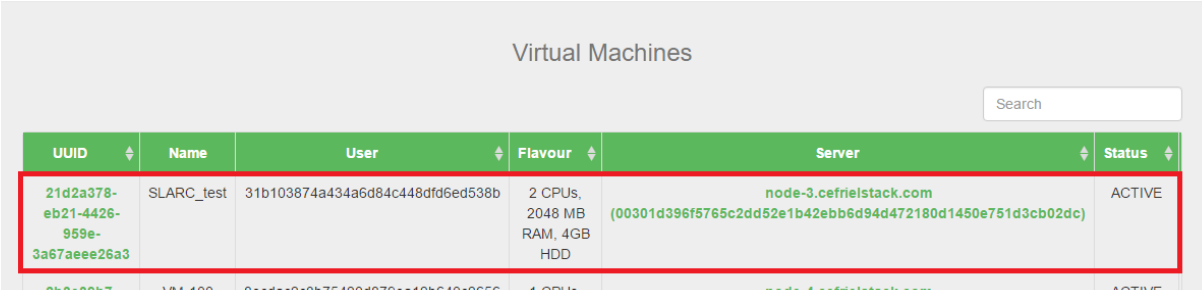

Figure 4-7: Initial VMs state for test 4.3.2. The VM of interest is highlighted in red.

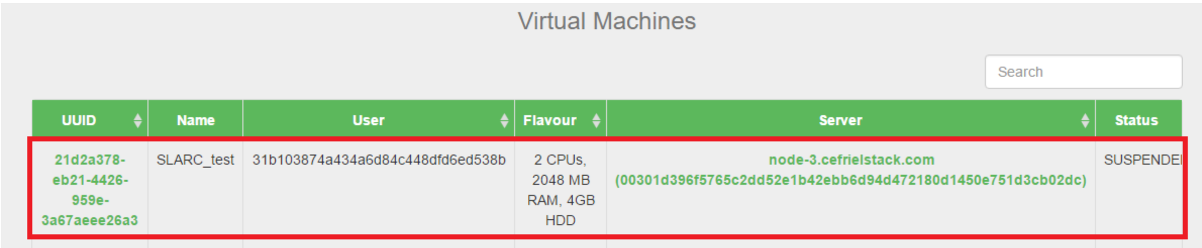

Figure 4-8: After issuing the pause command though the Policy Actuator, the status changed to SUSPENDED.

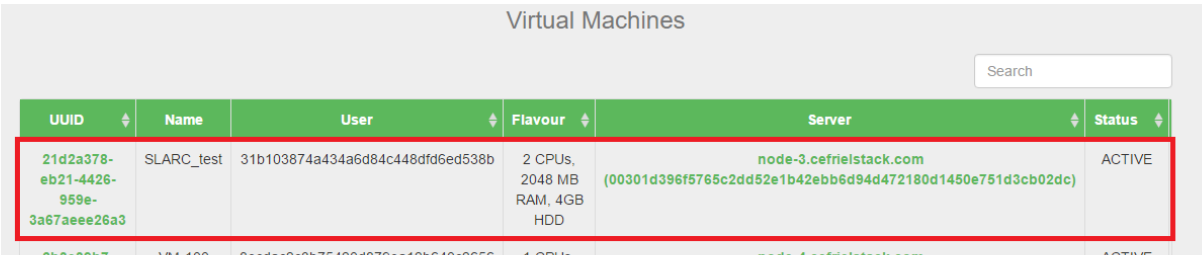

Figure 4-9: After executing the optimization plan, the VM is again in ACTIVE state.

The relevant actions that triggered the above-presented VM state changes were logged into the eCOP Monitor Database (the difference in the time of the logged actions are due to the fact that the User Interface of the eCOP presents the time in local context, rather than UTC).

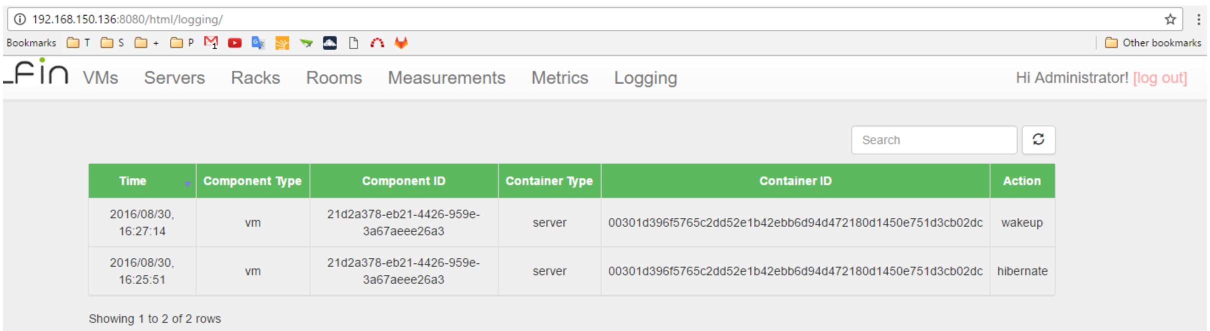

Figure 4-10: The DOLFIN eCOP Monitor DB has logged the relevant actions performed by the Policy Actuator.

4.3.2.6 Testing scenario check points

4.3.2.6.1 A VM gets into SUSPENDED state to emulate a forced shutdown is issued through the Policy Actuator

Relevant evidence is provided by means of Figure 4-7, Figure 4-8 and Figure 4-10. The VM was initially in ACTIVE state, it was changed to SUSPENDED and the relevant action has been logged into the eCOP Monitor Database logging interface.

Considering the above, this checkpoint was verified.

4.3.2.6.2 The Policy Actuator notifies the SLARC of the change in the VM state

As per Table 4-7, this step is to be validated based on the SLARC logs. Indeed, when checking the logs, the following lines were identified:

```
Message received: {"payload":{"name":"null-21d2a378-eb21-4426-959e-
3a67aeee26a3","eventType":"STOP_VM_CRASH","id":"21d2a378-eb21-4426-959e-
3a67aeee26a3"},"message":"event_notification","timestamp":"1472563560632"} On
Topic: vim With routingKey: key.manager and Content type: text/plain
```

Indicating that the event was received at Tue, 30 Aug 2016 13:26:00.632 UTC time.

Considering the above, this checkpoint was verified.

4.3.2.6.3 After 60 seconds, the SLARC emits an alarm that the VM SLA is about to be broken

As documented in 4.3.2.3, SLARC was set to emit a VM availability alarm after detecting 60 seconds of entering a non-ACTIVE state. Indeed, looking at the logs, one could find the line:

```
13:27:00,591 INFO com.espertech.esper.Timer-default-0 AmqpPublisher:publish:47 -
Message published: {"message":"Event
warning","timestamp":1472563620585,"payload":{"vmId":"21d2a378-eb21-4426-959e-
3a67aeee26a3","hostId":" ","type":"AVAILABILITY
VIOLATION","level":"CRITICAL","timeRepresentation":"duration","time":60,"penalty":0
.5}}__on exchange(topic): aggregationResult_AND RKey: aggregationResult
```

The above log entry indicates that SLARC acted after 59.953 msec of VM inactivity.

Considering the above, this checkpoint was verified.

4.3.2.6.4 The Policy Maker changes the DC Policy to optimize the SLA of the VM in hand and notifies the Optimizer

When an SLA-related alarm gets triggered, the Policy Maker should receive this alarm and properly handle it. The logs of the Policy Maker contained the following entries:

```
SlaHandler received b'{"message":"Event
warning","timestamp":1472563620585,"payload":{"vmId":"21d2a378-eb21-4426-959e-
3a67aeee26a3","hostId":" ","type":"AVAILABILITY
VIOLATION","level":"CRITICAL","timeRepresentation":"duration","time":60,"penalty":0
.5}}'

SLA violation with VM id 21d2a378-eb21-4426-959e-3a67aeee26a3
```

```
Sending violation alert with uuid 21d2a378-eb21-4426-959e-3a67aeee26a3
```

The above indicate that the alarm was successfully captured. Next, the following log lines indicate that the Policy Maker

1. Retrieved the available Policies and constraints from the Policy Repository;
2. Formulated a proper DC Policy to send it to the optimizer;
3. Authenticated against the Optimizer;
4. Set the active Policy for the optimization.

1	`HTTP GET http://127.0.0.1:8084/v1/policies/`
2	`Sending message to optimizer: {'constraints': [{'name': 'preserve_vm_performances', 'value_type': 'boolean', 'descr': 'Prevent engaging in actions which may degrade actual VM performances (observe SLA performance limits)'}, {'name': 'do_not_stop_vms', 'value_type': 'boolean', 'descr': 'Prevent VM shutdowns or suspensions'}, {'name': 'avoid_standbys', 'value_type': 'boolean', 'descr': 'Avoid noticeable (SLA infringing) VM standbys'}, {'name': 'do_not_accept_xdc_inbound_migrations', 'value_type': 'boolean', 'descr': 'Reject any and all inbound cross-DC workload migration requests'}, {'value': '21d2a378-eb21-4426-959e-3a67aeee26a3', 'name': 'SLA_violation_uuid', 'value_type': 'number', 'descr': 'Virtual machine uuid which violates the SLA Agreement'}], 'is_enforced': False, 'is_active': True, 'descr': 'attempt to reslve SLA violations', 'target': 'sla', 'name': 'sla_rebalance', 'conflicts_with': []}`
3	`HTTP POST http://127.0.0.1:8086/api/tokens/`
4	`HTTP POST http://127.0.0.1:8086/api/policies/active/`

After these steps, the Optimizer was featuring an updated Policy:


Details of Optimization policy 1 (sla_rebalance) with target Sla.

Status: The policy is currently active
Description: Attempt To Reslve Sla Violations.

Related Constraints

Name	Description	Value Type	Value
avoid_standbys	Avoid Noticeable (Sla Infringing) Vm Standbys	Boolean	None
do_not_accept_xdc_inbound_migrations	Reject Any And All Inbound Cross-Dc Workload Migration Requests	Boolean	None
do_not_stop_vms	Prevent Vm Shutdowns Or Suspensions	Boolean	None
preserve_vm_performances	Prevent Engaging In Actions Which May Degrade Actual Vm Performances (Observe Sla Performance Limits)	Boolean	None
SLA_violation_uuid	Virtual Machine Uuid Which Violates The Sla Agreement	Number	21d2a378-eb21-4426-959e-3a67aeee26a3

Figure 4-11: The active policy of the DC, as perceived by the Optimizer during test 4.3.2.

Considering the above, this checkpoint was verified.

4.3.2.6.5 The Policy Maker issues an optimization request to the Optimizer

Continuing the log inspection, it was identified that a new Optimization request was issued to the Optimizer:

```
HTTP POST http://127.0.0.1:8086/api/requests/
```

The User Interface of the Optimizer reported[2]:

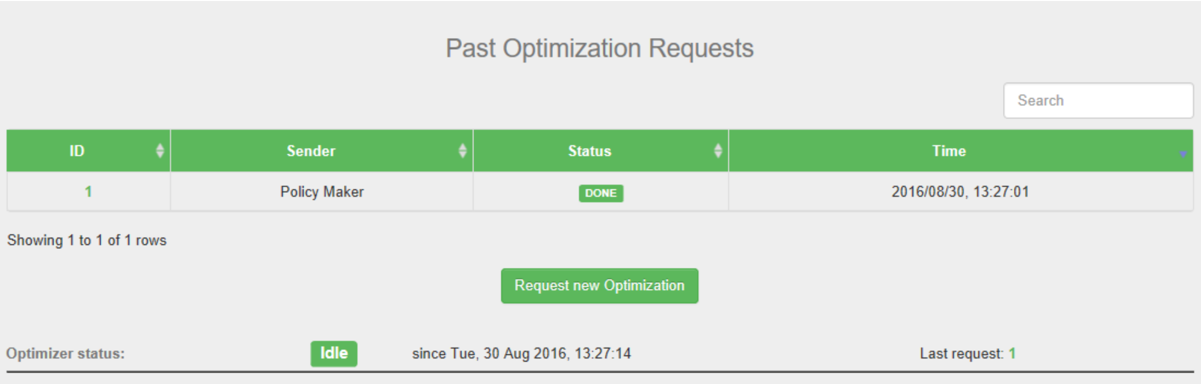

Figure 4-12: The optimization request issued by the Policy Maker, in the context of test 4.3.2.

Considering the above, this checkpoint was verified.

4.3.2.6.6 The Optimizer schedules a new optimization to wake the VM in hand up

After the receipt of the request, the Optimizer indeed started an optimization request, as depicted in Figure 4-13.

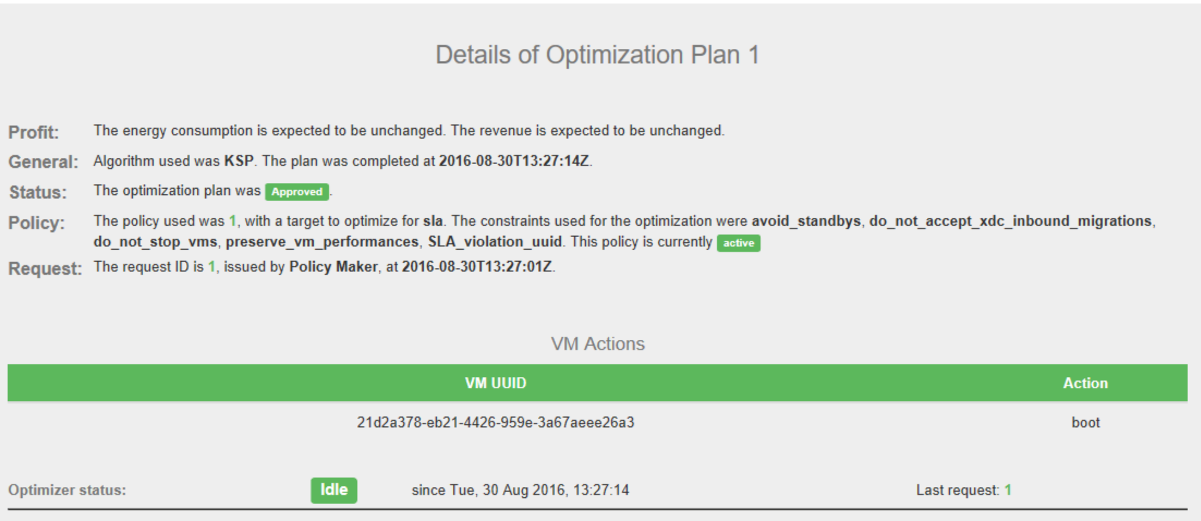

Figure 4-13: The optimization plan devised by the Policy Maker, in the context of test 4.3.2.

[2] The screenshot refers to a time instance after the optimization determination process. The times presented in the Optimizer UI are in local time (EEST at the point).

As can be observed from Figure 4-13, the Policy used was the one with ID 1 (the same as in Figure 4-11) and the request ID was also 1 (the same as in Figure 4-12), indicating that the depicted plan was generated as a response to the Policy Maker actuation against the SLA alarm emitted by SLARC. Also, the only action that the optimization plan in hand contains was related to booting the VM that was in SUSPENDED state.

Considering the above, this checkpoint was verified.

4.3.2.6.7 The Optimizer requests approval from the Policy Maker to forward the plan to the Policy Actuator

As can be seen from Figure 4-13, the optimization plan has an "Approved" status, indicating that the Policy Maker approved the plan. This can also be verified from the Policy Maker logs:

```
received optimizer request id: 1

POST /v1/policy/optimizer/approve (127.0.0.1) 1.21ms
```

Considering the above, this checkpoint was verified.

4.3.2.6.8 The Optimizer forwards the plan to the Policy Actuator

Evidence of this interaction can be found in the Policy Actuator logs:

1	2016-08-30 13:27:14 INFO PolicyActuator:214 - Received request to wakeupVMs { "vms" : [{ "vm_uuid" : "21d2a378-eb21-4426-959e-3a67aeee26a3", "start_time" : "2016-08-30T12:27:14Z" }] }
2	2016-08-30 13:27:14 INFO PolicyActuator:229 - Logged to eCOP DB { "component_type" : "vm", "component_id" : "21d2a378-eb21-4426-959e-3a67aeee26a3", "container_type" : "server", "container_id" : "-", "action" : "wakeup" }
3	2016-08-30 13:27:14 INFO PolicyActuator:238 - Delegating to adapter action PolicyActuatorAction [id=10, time=Tue Aug 30 13:27:14 UTC 2016, action=WAKEUP, componentType=VM, componentId=21d2a378-eb21-4426-959e-3a67aeee26a3, actionParameters=[21d2a378-eb21-4426-959e-3a67aeee26a3, Tue Aug 30 12:27:14 UTC 2016], actionParameterClasses=[class java.lang.String, class java.util.Date], actionParameterNames=[vm_uuid, start_time]] 2016-08-30 13:27:14 INFO ComputeActions:238 - Resume Server
4	2016-08-30 13:27:16 INFO ComputeActions:243 - Server 21d2a378-eb21-4426-959e-3a67aeee26a3 resumed.
5	2016-08-30 13:27:16 INFO AmqpPublisher:47 - Message published: {"payload":{"name":"null-21d2a378-eb21-4426-959e-3a67aeee26a3","eventType":"START_VM","id":"21d2a378-eb21-4426-959e-3a67aeee26a3"},"message":"event_notification","timestamp":"1472563636821"}__on exchange(topic): vim_AND RKey: key.manager

The above indicate that the Policy Actuator successively:

1. Received the VM-wakeup directive from the Optimizer;

2. Logged the relevant action to the eCOP DB;
3. Generated a command to the underlying DCO Hypervisor Broker to wake the VM up;
4. Acknowledged the VM wakeup;
5. Informed the SLARC about the VM state change.

Considering the above, this checkpoint was verified.

4.3.2.6.9 The Policy Actuator wakes the VM up

As per Figure 4-9, it can be deduced that the VM is ACTIVE and therefore was woken up.

Hence, this checkpoint was verified.

4.3.3 Smart Grid testing scenario

DOLFIN is designed to monitor a large number of metrics, assisting towards an environmental friendly computing infrastructure. To this end, efficient use of power provided from the Smart Grid is a key to achieve this objective. Interfacing with the Smart Grid so as to receive relevant information from the DC energy providers enables DOLFIN to increase power savings while at the same time reducing costs, as DOLFIN can take measures in response to Smart Grid notifications, such as price variations, renewable mix information and Grid usage statistics.

In this scenario we envisage to test DOLFIN behaviour against a series of common Smart Grid events and measure relevant KPIs that highlight the increased efficiency of DOLFIN-enabled DCs. In DOLFIN we integrate the Open Automated Demand Response (OpenADR) [27] protocol to test our approach. OpenADR is a communications protocol designed to facilitate transmission and reception of demand-response signals from a utility or independent system operator to electricity customers. The DOLFIN Smart Grid Controller (SGC) acts as a gateway, receiving demand-response events from a Demand Response Automation Server and converting them into information for the Energy Efficiency Policy Maker and Actuator. The basic scenario description is the following:

The ADR Server will provide to the SGC daily information about the energy prices in one hour intervals. At random points during the day, the electricity prices will be adjusted to reflect a new situation in the grid generation facilities and the ADR Server will inform the SGC of these changes. The SGC will, next, store the values to the DOLFIN Information DB. The DCO

4.3.3.1 Requirements addressed by the testing scenario

The main objective of this testing scenario is to show DOLFIN's responsiveness under extreme DR requests (steep price change) from the Smart Grid side.

The DOLFIN components under examination in this scenario are:

- Smart Grid Controller;
- ICT Performance & Energy Supervisor;

- Energy efficiency policy maker and actuator;
- eCOP Monitor DB;
- DCO Brokers;
- Cross DC Workload Orchestrator;

4.3.3.2 Test prerequisites

To successfully handle the test scenario, the following prerequisites are needed:

1. Two physically separated (different chassis, different networks, different power source) cloud computing environments with active VMs with intensive CPU load. The DC resources are managed by DOLFIN, which is supported by OpenStack or any other OpenStack API-compatible installation.
2. A DOLFIN instantiation comprising all relevant components identified as test components.
3. Proper monitoring equipment should have been deployed to monitor the performance and characteristics of the DC elements of interest, including physical servers, server racks, power supplies.
4. The components identified in 4.3.3.2 are properly setup and configured

4.3.3.3 Testing setup and configuration

In order to perform the tests under controlled conditions, the OpenADR server will be deactivated for the test. Instead, we will emulate the OpenADR server by emitting the signals it would emit if it were operating normally. On the other hand, the SGC will be listening for such messages in order to be aware of any price changes. The SGC will build an internal calendar of prices, and send the price data on a pre-defined output interface in the DOLFIN Information DB, to be persisted.

The two distinct DCs are represented by two separate OpenStack installations (Figure 4-14). Each installation represents different OpenStack Region. Region One is a virtualized Mirantis [28] OpenStack distribution deployed and managed by Openstack Fuel [29]. RegionOne server is supplied by renewable energy generated by photovoltaics system. RegionTwo is a standard OpenStack installation prepared according to the official guides [24] [30]. Both regions share the same authorization service Keystone. Connection between both regions is provided by 25km optic fibre that represents real distance between two DCs. The two regions are running on different racks and contain different servers, hosting, among other VMs, two DOLFIN instances (one in RegionTwo and one in RegionOne). The following tables tabulate the details of the DC testbed.

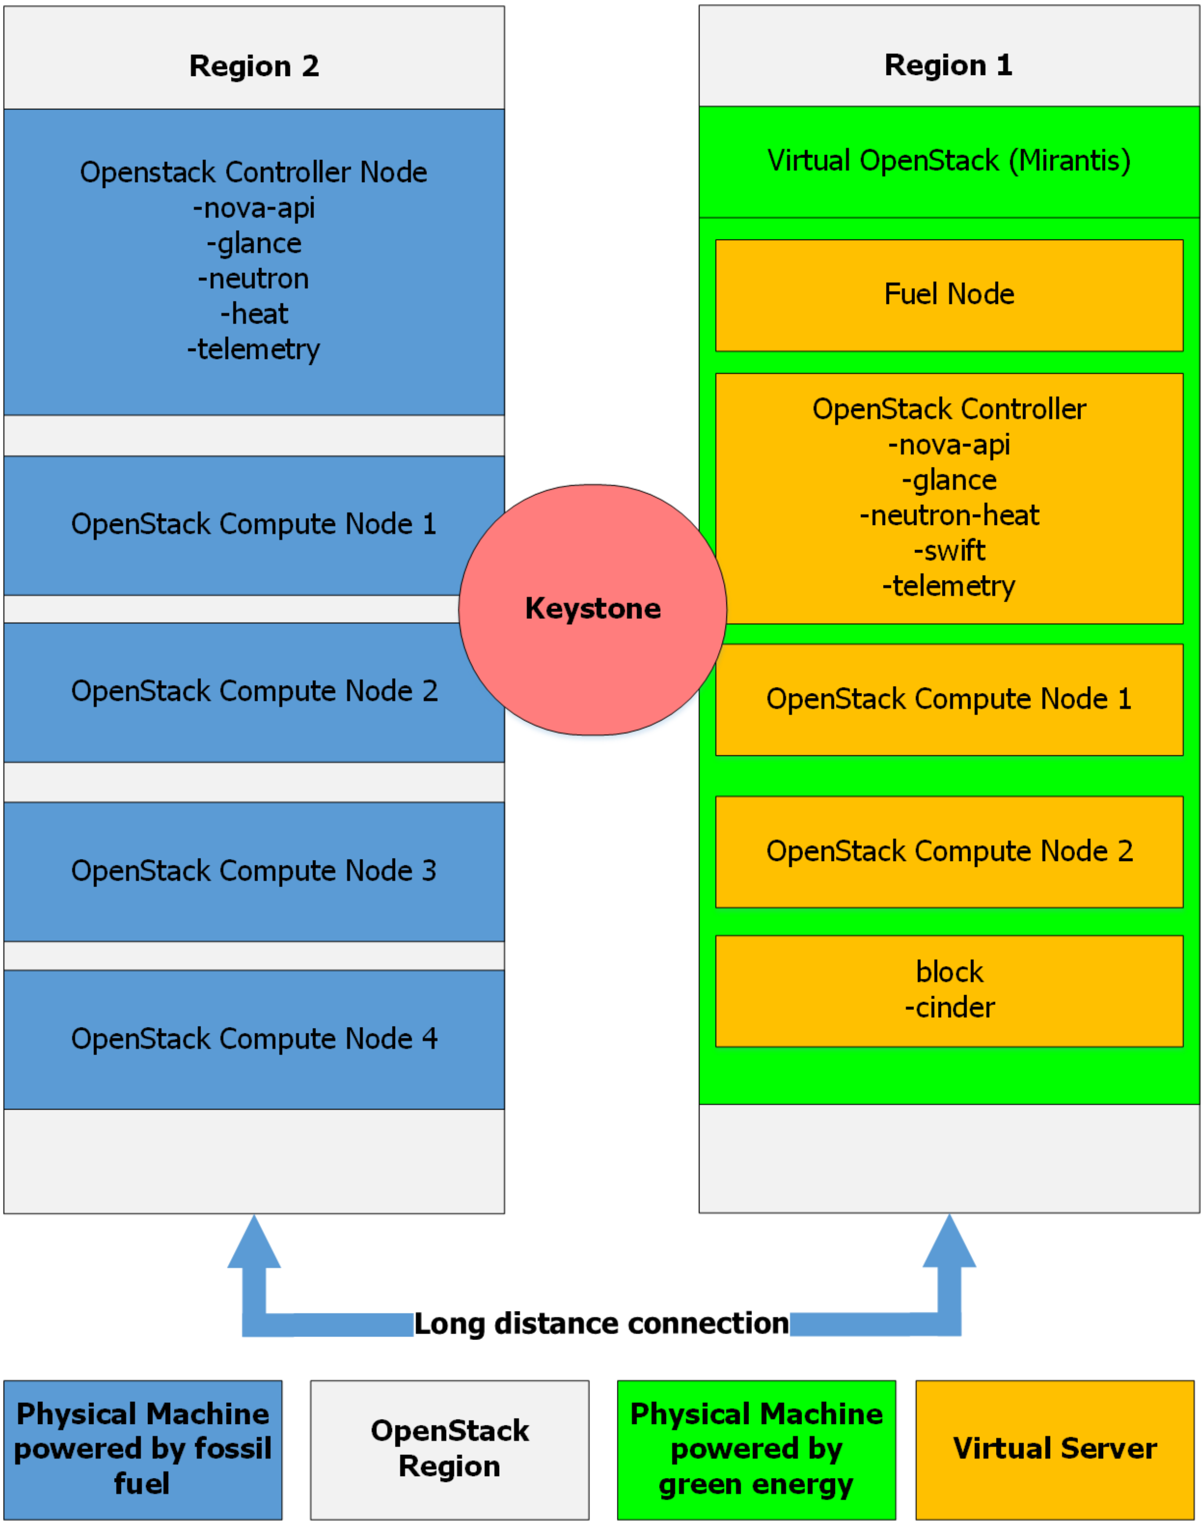

Figure 4-14 Smart Grid testing scenario's testbed architecture

Table 4-8: Physical servers (compute nodes) hosting the VMs in DC #1 (RegionTwo)

Serial Number	03d914e78243	8dfe147710a6	9f722d91d955	e1f2f114ee765

(Openstack Hypervisor ID)	**77b0542eeba56c8304793ae95d739e44d0e04ba7154c**	**3de848b9f48ef35df969636c70426bf3f432dc2249b2**	**4db9f06c1660d9607a92f6077dba0d038bc22f7e32b5**	**508569dec5a852f844f7c97ec2de4b32be531891ad5**
Name	air-5	air-6	compute4	compute3
Rack_id	1	1	1	1
Active	True	True	True	True
CPU	32	32	32	32
Ram (MB)	64038	64038	64037	64037
Hdd	853	853	853	853
Cpu_frequency	2400	2400	2400	2400
Min_watt_per_cpu_core	2.88	2.88	2.88	2.88
Max_watt_per_cpu_core	6.00	6.00	6.00	6.00
Min_watt_per_kbps	0.00	0.00	0.00	0.00
Max_watt_per_kbps	0.00	0.00	0.00	0.00
Cpu_energy_mult	0.80	0.80	0.80	0.80
Ram_energy_mult	0.20	0.20	0.20	0.20
Net_energy_mult	0.00	0.00	0.00	0.00
Green	False	False	False	False

Note that all servers are cooled by brown energy sources (energy coming from the grid) and are air-cooled. Next, the server emulating the second DC (through 2-layer virtualization) is presented tabulated. Note that this server is cooled with liquid cooling and that it is powered by local green energy sources (photovoltaics).

Table 4-9: Physical node of DC #2 (RegionOne)

Serial Number (Openstack Hypervisor ID)	**03d914e7824377b0542eeba56c8304793ae95d739e44d0e04ba7154c**
Name	Xeon_18
Rack_id	1
Active	True
CPU	32
Ram (MB)	64038
Hdd	853
Cpu_frequency	2400
Min_watt_per_cpu_core	2.88
Max_watt_per_cpu_core	6.00
Min_watt_per_kbps	0.00
Max_watt_per_kbps	0.00
Cpu_energy_mult	0.80
Ram_energy_mult	0.20
Net_energy_mult	0.00
Green	True

The above servers host a total of 66 VMs, as tabulated below (the VMs have been grouped based on their host servers and characteristics for reasons of brevity):

Table 4-10: VMs hosted in DC #1 (RegionTwo)

# of VMs	Compute Node	VCPUs	VRAM (MB)
2	air-5	1	128
1	air-5	2	128
6	air-5	1	1024
27	air-5	2	1024
18	air-5	4	1024
1	air-5	4	4096
5	air-6	1	1024
3	air-6	2	1024
2	air-6	4	1024

Table 4-11: VMs hosted in DC #2 (RegionOne)

# of VMs	Compute Node	VCPUs	VRAM (MB)
1	Xeon_18	4	4096

As far as networking is concerned, both RegionOne and RegionTwo feature 10Gb interfaces. Although the test scenario was carried out in laboratory conditions, in order to simulate more realistic circumstances length of the fibre connection was artificially prolonged up to 25 kilometres (maximum fibre length supported by the interfaces). This allowed us to add more delay between the two regions.

In the case of VM live migration (or even entire DCs) large volume of data is going to be copied over the network. Therefore it is important to have an up-to-date view of interconnections between Data Centers in order to ensure as quickly as possible data transfers. DOLFIN platform uses Cross-DC Network Monitoring component to constantly measure bandwidth values between distributed DCs. Underneath the module runs a set of iPerf instances that collects the results. The following throughput performance was recorded between both DCs (Openstack Regions: RegionOne and RegionTwo):

Table 4-12: The testbed communication channel details.

Interval[s]	Transfer[GBytes]	Bandwidth[Gbits/sec]
0.0-1.00	1.10	9.41
1.00-2.00	1.09	9.41

2.00-3.00	1.10	9.42
3.00-4.00	1.10	9.42
4.00-5.00	1.09	9.41
5.00-6.00	1.10	9.42
6.00-7.00	1.09	9.41
7.00-8.00	1.10	9.42
8.00-9.00	1.09	9.41
9.00-10.00	1.10	9.42

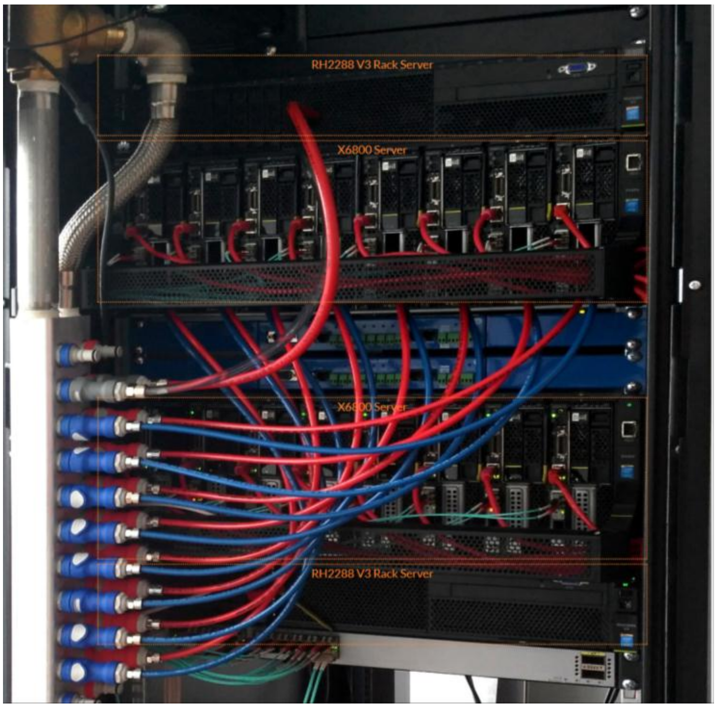

Figure 4-15: The rack hosting the DC #1 physical nodes.

It should be underlined that according to the measurements acquired by the DCO Brokers, the average CPU utilization of the RegionTwo VMs, in total, was around 18%:

```
$ mysql -uroot -p -e "select AVG(value) from ecop.generic_measurements where
type_id='cpu_util' and resource='vm' and time between '2016-09-11 10:00:00' and
```

```
'2016-09-11 11:00:00'"

Enter password:
+------------+
| AVG(value) |
+------------+
| 17.899325  |
+------------+
```

4.3.3.4 Test execution and expected results

Initially, the energy price for the day will be set to a flat 0.25 EUR/kWh by sending an OpenADR message to set the price in an absolute manner. In fact, the prices need not to be flat, but will be set likewise for reasons of simplicity. Next, a message to change the prices relatively to the first one will be sent, increasing the price of the energy by 0.15 EUR/kWh, reaching a total of 0.40 EUR/kWh, indicating that the Smart Grid Operator wishes to reduce the energy consumption of the Grid as a whole. The Policy Maker, perceiving the significant price change (60%) will, then, notify the Optimizer to perform an optimization plan by keeping the DC energy expenses as close as possible to the ones before the price change, capping the energy dissipation of the DC at 62.5%. The next table presents the expected test execution steps:

Step #	Test Action	Expected Results	Means of Validation
1.	A message to statically set the prices to 0.25EUR/kWh is sent to the SGC	The SGC will consume the message and build an internal calendar, also logging the new price value in the DOLFIN Info DB.	The new prices will be available for retrieval from the DOLFIN Info DB.
2.	A message to increase the energy price by 60% is sent to the SGC	The SGC will consume the message and build an internal calendar, also logging the new price value in the DOLFIN Info DB.	The new prices will be available for retrieval from the DOLFIN Info DB.
3.	The Policy Maker detects the severe price change and changes the DC Policy also setting it to the Optimizer	The Optimizer changes its policy to cap the DC Energy consumption to 62.5%	The new policy of the optimizer is set to cap the energy consumption to 62.5%, allowing for XDC migrations
4.	The Policy Maker notifies the optimizer to generate a new optimization plan	The Optimizer receives the notification and initiates the optimization procedure.	A new optimization plan is generated by the Optimizer, respecting the policy above.

5.	The Optimizer sends the generated plan to the Policy Maker to approve it	The Policy Maker contacts the Cross DC Workload Orchestrator who responds that a candidate DC has been found (RegionOne) and, then, approves the plan	The generated optimization plan should appear as "approved"
6.	The Optimizer forwards the plan to the Policy Actuator to implement the VM migrations and the server actions and to the Policy Maker to organize the VMs relocation	The VM and server actions get actuated by the Policy Actuator and the VM relocations are forwarded to the Cross DC VM Manager by the Policy Maker	Evidence of the actions initiation should be found in the eCOP DB Broker logging facility, in the logs of the components and in the DOLFIN Info DB (for the VM relocations)
7.	Check that the energy consumption of the DC has fallen at or under the requested limit.	The energy consumption of the DC should have decreased to at most 62.5%.	The eCOP DB should provide evidence of the change in the DC energy consumption.

Table 4-13 – Test actions for testing scenario 3

On test completion, the testing scenario should demonstrate:

- The price changes are properly handled by the SGC and that the Policy Maker successfully detects and handles the price changes.
- The VM migrations, server actions, and VM relocations are properly handled by the Policy Actuator and the Cross DC VM Manager, respectively.
- The original trigger condition is met (e.g. to achieve at least 37.5% of energy reduction);

4.3.3.5 Outcome of the test

After the test was completed, we could see that the Optimizer had indeed generated an optimization plan that, after executed, resulted in reducing the energy consumption of the DC, setting it to 54.46% of the original energy consumption (energy consumption of the rack representing the DC after the application of the optimization plan was approximately 310W, starting from approximately 560W, before the application of the optimization plan). As the initially requested change (from the Policy Maker) was to cap the DC energy consumption at 62.5% of the initial one, the test is considered successful.

It should be highlighted that the plan generated by the Optimizer was considered optimal; after the optimization plan execution, only one server remained active (namely air-5), its RAM being almost depleted due to VMs operation[3]. If the Optimizer would like to host one more VM in DC #1 (RegionTwo), then another server should be made active, resulting to an additional energy expense of 92W (32 cores x 2.88 W/core in idle state). In this case, the energy consumption of the DC would rise to 402W, being capped at approximately 71% of the initial energy consumption, failing to respect the request issued by the Policy Maker to reduce the energy consumption down to 62.5% of the initial.

Figure 4-166 depicts the energy consumption of the server before and after the implementation of the optimization plan.

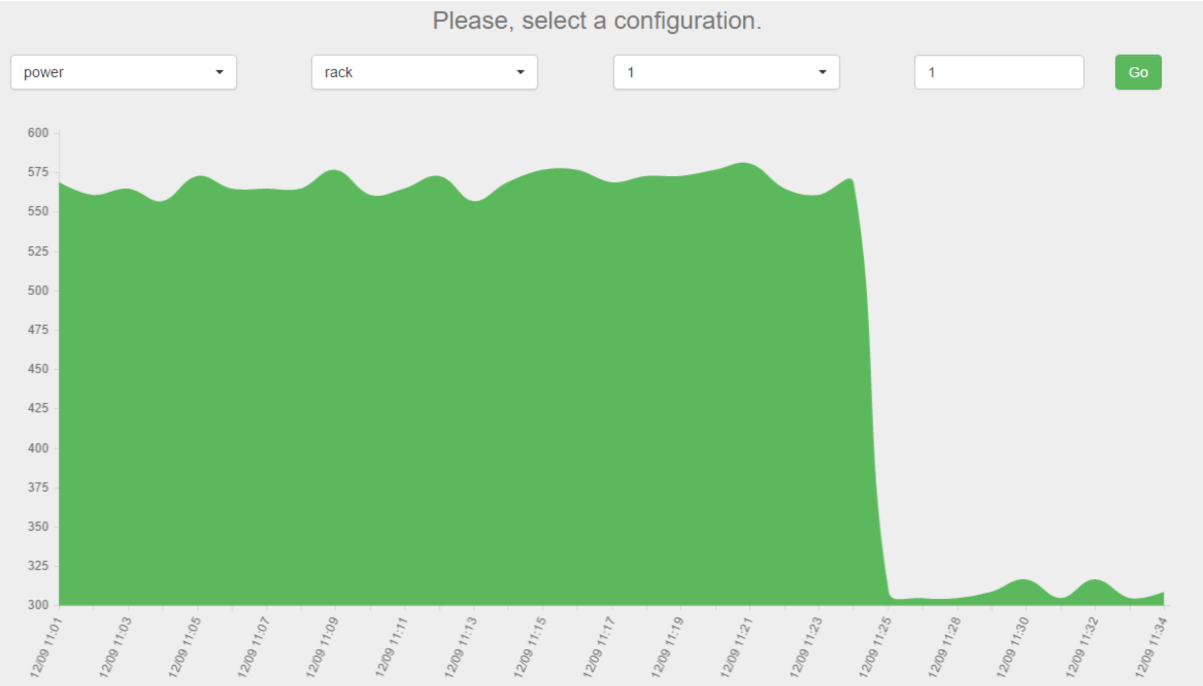

Figure 4-16: Energy consumption of the rack representing DC #1 (RegionTwo) before and after the implementation of the optimization plan.

4.3.3.6 Testing scenario check points

4.3.3.6.1 A message to statically set the prices to 0.25EUR/kWh is sent to the SGC

As already stated, for reasons of predictability, we bypassed the ADR server in the rendering of the OpenADR events, feeding the SGC directly with the prices we desired for performing the test. To set the price statically to 0.25EUR/kWh, the following message was sent to the RabbitMQ interface normally exposed by the ADR component, which is also listened by the SGC.


```
{"endDate":"2016-09-11        23:59","name":"Daily        energy        prices
[EUR/kW]","responseRequired":false,"schedule":[{"interval":{"end":"00:59","start":"
00:00"},"value":0.25},{"interval":{"end":"01:59","start":"01:00"},"value":0.25},{"i
```


[3] By default, Openstack sets memory over commissioning to 1, namely each VM has a dedicated amount of memory as dictated by its flavour. In this framework, when the server's memory is full, no more VMs can be accepted for hosting.

nterval":{"end":"02:59","start":"02:00"},"value":0.25},{"interval":{"end":"03:59","start":"03:00"},"value":0.25},{"interval":{"end":"04:59","start":"04:00"},"value":0.25},{"interval":{"end":"05:59","start":"05:00"},"value":0.25},{"interval":{"end":"06:59","start":"06:00"},"value":0.25},{"interval":{"end":"07:59","start":"07:00"},"value":0.25},{"interval":{"end":"08:59","start":"08:00"},"value":0.25},{"interval":{"end":"09:59","start":"09:00"},"value":0.25},{"interval":{"end":"10:59","start":"10:00"},"value":0.25},{"interval":{"end":"11:59","start":"11:00"},"value":0.25},{"interval":{"end":"12:59","start":"12:00"},"value":0.25},{"interval":{"end":"13:59","start":"13:00"},"value":0.25},{"interval":{"end":"14:59","start":"14:00"},"value":0.25},{"interval":{"end":"15:59","start":"15:00"},"value":0.25},{"interval":{"end":"16:59","start":"16:00"},"value":0.25},{"interval":{"end":"17:59","start":"17:00"},"value":0.25},{"interval":{"end":"18:59","start":"18:00"},"value":0.25},{"interval":{"end":"19:59","start":"19:00"},"value":0.25},{"interval":{"end":"20:59","start":"20:00"},"value":0.25},{"interval":{"end":"21:59","start":"21:00"},"value":0.25},{"interval":{"end":"22:59","start":"22:00"},"value":0.25},{"interval":{"end":"23:59","start":"23:00"},"value":0.25}],"scheduleType":"**STATIC**","startDate":"2016-09-11 00:00","timestamp":1473592396291,"typeID":"**PRICE_ABSOLUTE**"}

Then, we noticed the following log lines in the SGC logs:

1	MultiplePriceNotifier: price at [2016/09/11 11:13:16.293] -> 0.25 MultiplePriceNotifier: price at [2016/09/11 12:13:16.293] -> 0.25 MultiplePriceNotifier: price at [2016/09/11 13:13:16.293] -> 0.25 MultiplePriceNotifier: price at [2016/09/11 14:13:16.293] -> 0.25 MultiplePriceNotifier: price at [2016/09/11 15:13:16.293] -> 0.25 MultiplePriceNotifier: price at [2016/09/11 16:13:16.293] -> 0.25
2	MultiplePriceNotifier: calling http://localhost:8090/prices {"message":"energy_price_window","payload":{"prices":[0.25,0.25,0.25,0.25,0.25,0.25],"type":"PRICE_ABSOLUTE"},"timestamp":1473592396293}

Namely, SGC correctly identified the statically set absolute energy price (1st logged line) and, next, it notified the DOLFIN Info DB about the price change.

Considering the above, this checkpoint was verified.

4.3.3.6.2 A message to increase the energy price by 60% is sent to the SGC

The message was communicated to the SGC via the dedicated RabbitMQ interface normally exposed by the ADR component, which is also listened by the SGC. The message sent was instructing the SGC to consider that the price was (relatively) increased by 0.15 EUR/kWh and was as follows:

{"endDate":"2016-09-11 23:59","name":"Daily energy prices [EUR/kW]","responseRequired":false,"schedule":[{"interval":{"end":"00:59","start":"00:00"},"value":0.15},{"interval":{"end":"01:59","start":"01:00"},"value":0.15},{"interval":{"end":"02:59","start":"02:00"},"value":0.15},{"interval":{"end":"03:59","start":"03:00"},"value":0.15},{"interval":{"end":"04:59","start":"04:00"},"value":0.15},{"interval":{"end":"05:59","start":"05:00"},"value":0.15},{"interval":{"end":"06:59","start":"06:00"},"value":0.15},{"interval":{"end":"07:59","start":"07:00"},"value":0.15},{"interval":{"end":"08:59","start":"08:00"},"value":0.15},{"interval":{"end":"09:59","start":"09:00"},"value":0.15},{"interval":{"end":"10:59","start":"10:00"},"value":0.15},{"interval":{"end":"11:59","start":"11:00"},"value":0.15},{"interval":{"end":"12:59","start":"12:00"},"value":0.15},{"interval":{"end":"13:59","start":"13:00"},"value":0.15},{"interval":{"end":"14:59","start":"14:00"},"value":0.15},{"interval":{"end":"15:59","start":"15:00"},"value":0.15},{"interval":{"end":"16:59","start":"16:00"},"value":0.15},{"interval":{"end":"17:59","start":"17:00"},"value":0.15},{"interval":{"end":"18:59","start":"18:00"},"value":0.15},{"interval":{"end":"19:59","start":"19:00"},"value":0.15},{"interval":{"end":"20:59","start":"20:00"},"value":0.15},{"interval":{"end":"21:59","start":"21:00"},"value":0.15},{"interval":{"end":"22:59","start":"22:00"},"value":0.15},{"interval":{"end":"23:59","start":"23:00"},"value":0.15}],"scheduleType":"**DYNAMIC**","startDate":"2016-09-11 00:00","timestamp":1473592398254,"typeID":"**PRICE_RELATIVE**"}

From the logs of the SGC, one could see the following activity:

1	Message received: topic: adr routingKey: event.price {"endDate":"2016-09-11 23:59","name":"Daily energy prices [EUR/kW]","responseRequired":false,"schedule":[{"interval":{"end":"00:59","start":"00:00"},"value":0.15},{"interval":{"end":"01:59","start":"01:00"},"value":0.15},{"interval":{"end":"02:59","start":"02:00"},"value":0.15},{"interval":{"end":"03:59","start":"03:00"},"value":0.15},{"interval":{"end":"04:59","start":"04:00"},"value":0.15},{"interval":{"end":"05:59","start":"05:00"},"value":0.15},{"interval":{"end":"06:59","start":"06:00"},"value":0.15},{"interval":{"end":"07:59","start":"07:00"},"value":0.15},{"interval":{"end":"08:59","start":"08:00"},"value":0.15},{"interval":{"end":"09:59","start":"09:00"},"value":0.15},{"interval":{"end":"10:59","start":"10:00"},"value":0.15},{"interval":{"end":"11:59","start":"11:00"},"value":0.15},{"interval":{"end":"12:59","start":"12:00"},"value":0.15},{"interval":{"end":"13:59","start":"13:00"},"value":0.15},{"interval":{"end":"14:59","start":"14:00"},"value":0.15},{"interval":{"end":"15:59","start":"15:00"},"value":0.15},{"interval":{"end":"16:59","start":"16:00"},"value":0.15},{"interval":{"end":"17:59","start":"17:00"},"value":0.15},{"interval":{"end":"18:59","start":"18:00"},"value":0.15},{"interval":{"end":"19:59","start":"19:00"},"value":0.15},{"interval":{"end":"20:59","start":"20:00"},"value":0.15},{"interval":{"end":"21:59","start":"21:00"},"value":0.15},{"interval":{"end":"22:59","start":"22:00"},"value":0.15},{"interval":{"end":"23:59","start":"23:00"},"value":0.15}],"scheduleType":"DYNAMIC","startDate":"2016-09-11 00:00","timestamp":1473592398254,"typeID":"PRICE_RELATIVE"}
2	MultiplePriceNotifier: price at [2016/09/11 11:13:18.288] -> 0.4 MultiplePriceNotifier: price at [2016/09/11 12:13:18.288] -> 0.4 MultiplePriceNotifier: price at [2016/09/11 13:13:18.288] -> 0.4 MultiplePriceNotifier: price at [2016/09/11 14:13:18.288] -> 0.4 MultiplePriceNotifier: price at [2016/09/11 15:13:18.288] -> 0.4 MultiplePriceNotifier: price at [2016/09/11 16:13:18.288] -> 0.4
3	MultiplePriceNotifier: calling http://localhost:8090/prices {"message":"energy_price_window","payload":{"prices":[0.4,0.4,0.4,0.4,0.4,0.4],"type":"PRICE_ABSOLUTE"},"timestamp":1473592398288}

In the first log clip, evidence that the message originally sent by the ADR-emulating script was passed detected properly by the SGC. Next, the second log clip indicates that the SGC, aware of the previous energy price state, made the calculation in order come up with the updated energy price for the next six hours. Last, the third log clip indicates that this value was communicated to the DOLFIN Info DB.

Considering the above, this checkpoint was verified.

4.3.3.6.3 The Policy Maker detects the severe price change and changes the DC Policy also setting it to the Optimizer

When checking at the relevant Policy Maker logs, the following log clips could be detected:

1	2016-09-11 11:14:02,386 -> check_prices: main:80 Diff is 0.600000. Notifying the Optimizer
2	2016-09-11 11:14:05,044 -> check_prices: notify:45 Successuflly authenticated agaist the optimizer 2016-09-11 11:14:06,202 -> check_prices: notify:49 Successfully changed the policy {"is_active":true,"target":"energy","constraints":[{"type":"number","name":"energy_rel_value","value":62.5,"descr":"Maximum relative value for energy"},{"type":"boolean","name":"push_for_xdc_outbound_migrations","value":true,"descr":"Allow XDC load relocation"},{"type":"boolean","name":"do_not_stop_vms","value":true,"descr":"Do not STOP/PAUSE VMs"}]}

The first log snippet indicates that the price change (0.600000 corresponds to 60%) was detected by the Policy Maker. The second snippet indicates that the Policy Maker i) authenticated against the

Optimizer to get a valid token and ii) changed the policy of the Optimize, capping the overall DC consumption to 62.5% of the initial one. Figure 4-17, below, provides evidence of the policy set to the Optimizer via the Optimizer dashboard.

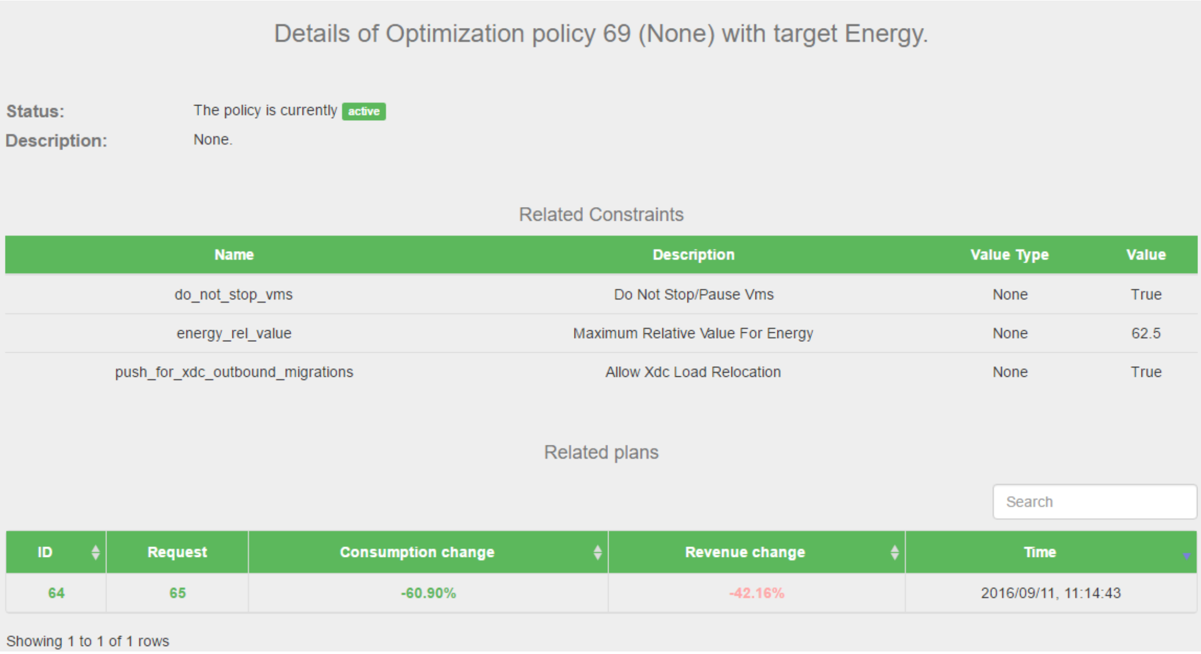
Details of Optimization policy 69 (None) with target Energy.

Status: The policy is currently active
Description: None.

Related Constraints

Name	Description	Value Type	Value
do_not_stop_vms	Do Not Stop/Pause Vms	None	True
energy_rel_value	Maximum Relative Value For Energy	None	62.5
push_for_xdc_outbound_migrations	Allow Xdc Load Relocation	None	True

Related plans

Search

ID	Request	Consumption change	Revenue change	Time
64	65	-60.90%	-42.16%	2016/09/11, 11:14:43

Showing 1 to 1 of 1 rows

Figure 4-17: The policy set to the Optimizer by the Policy Maker.

Considering the above, this checkpoint was verified.

4.3.3.6.4 The Policy Maker notifies the optimizer to generate a new optimization plan

Again looking at the Policy Maker logs, one could identify the following lines, according to which a request for optimization was sent by the Policy Maker at 2016-09-11 11:14 UTC.

```
2016-09-11 11:14:06,595 -> check_prices: notify:52
{"status":"RECV","sender":"Policy Maker","target":null,"id":65,"time":"2016-09-11T11:14:06.218944Z"}
```

This communication could be also checked through the dedicated Optimizer dashboard. Figure 4-18 provides the relevant evidence.

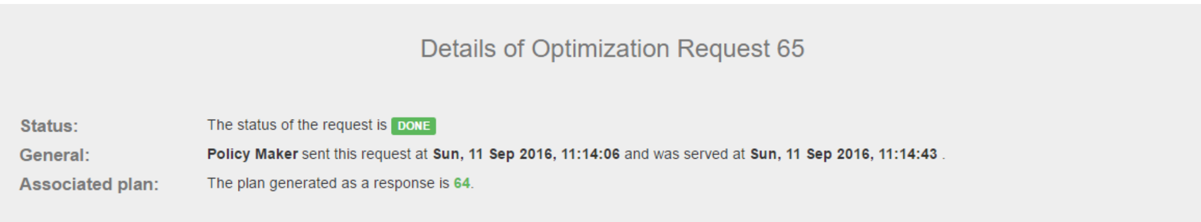
Details of Optimization Request 65

Status: The status of the request is DONE
General: **Policy Maker** sent this request at **Sun, 11 Sep 2016, 11:14:06** and was served at **Sun, 11 Sep 2016, 11:14:43** .
Associated plan: The plan generated as a response is 64.

Figure 4-18: The request performed by the Policy Maker, as perceived by the Optimizer.

As can be seen from Figure 4-18, in response to this request, an optimization plan with ID 64 was generated after 37 seconds. The details of the optimization plan are presented in the following figure.

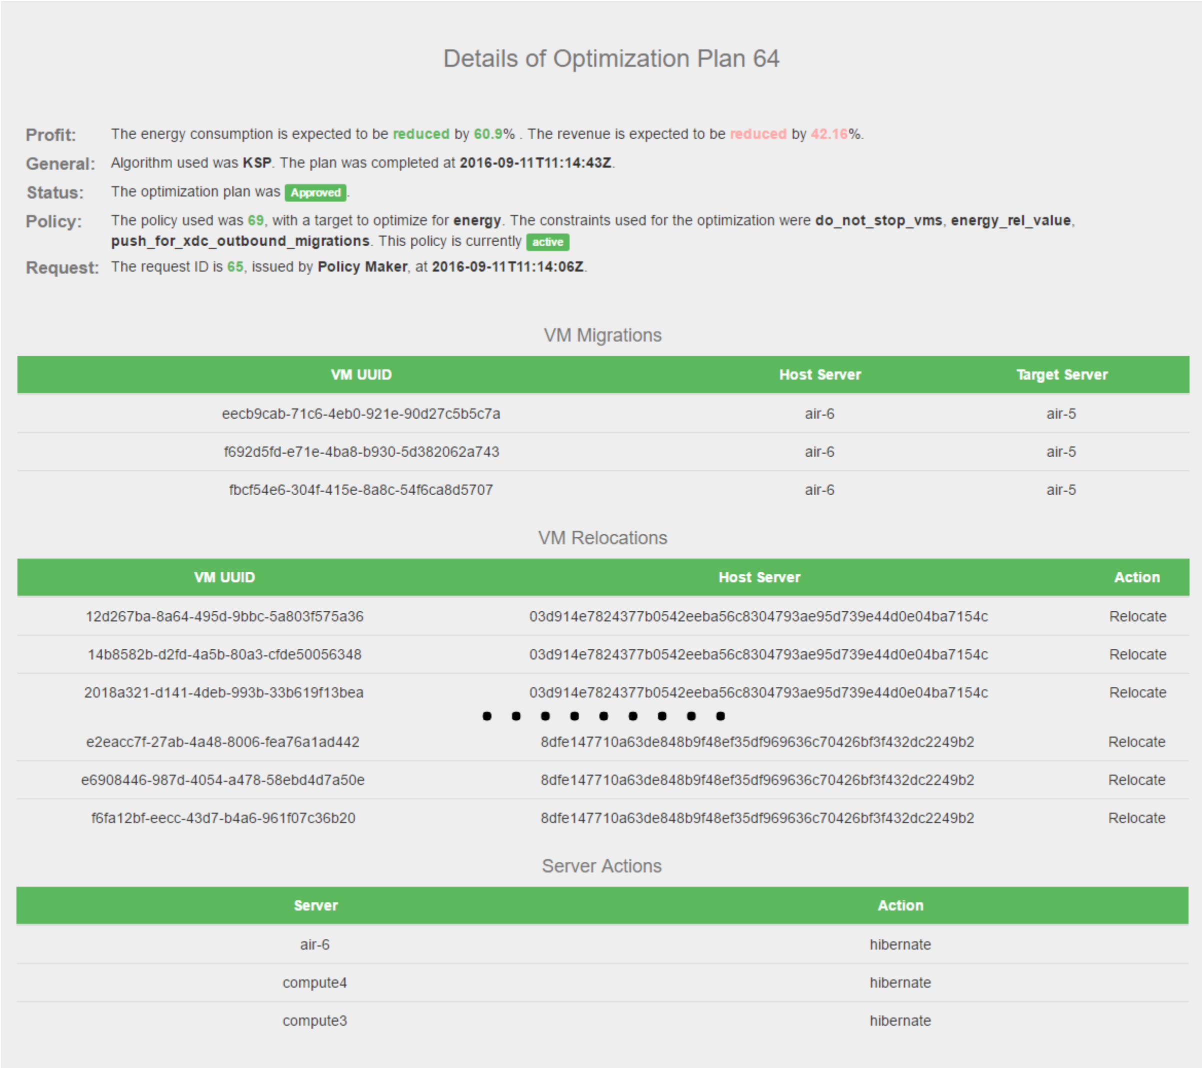

Details of Optimization Plan 64

Profit: The energy consumption is expected to be reduced by 60.9% . The revenue is expected to be reduced by 42.16%.
General: Algorithm used was **KSP**. The plan was completed at **2016-09-11T11:14:43Z**.
Status: The optimization plan was Approved.
Policy: The policy used was 69, with a target to optimize for **energy**. The constraints used for the optimization were **do_not_stop_vms**, **energy_rel_value**, **push_for_xdc_outbound_migrations**. This policy is currently active
Request: The request ID is 65, issued by **Policy Maker**, at **2016-09-11T11:14:06Z**.

VM Migrations

VM UUID	Host Server	Target Server
eecb9cab-71c6-4eb0-921e-90d27c5b5c7a	air-6	air-5
f692d5fd-e71e-4ba8-b930-5d382062a743	air-6	air-5
fbcf54e6-304f-415e-8a8c-54f6ca8d5707	air-6	air-5

VM Relocations

VM UUID	Host Server	Action
12d267ba-8a64-495d-9bbc-5a803f575a36	03d914e7824377b0542eeba56c8304793ae95d739e44d0e04ba7154c	Relocate
14b8582b-d2fd-4a5b-80a3-cfde50056348	03d914e7824377b0542eeba56c8304793ae95d739e44d0e04ba7154c	Relocate
2018a321-d141-4deb-993b-33b619f13bea	03d914e7824377b0542eeba56c8304793ae95d739e44d0e04ba7154c	Relocate
	• • • • • • • • •	
e2eacc7f-27ab-4a48-8006-fea76a1ad442	8dfe147710a63de848b9f48ef35df969636c70426bf3f432dc2249b2	Relocate
e6908446-987d-4054-a478-58ebd4d7a50e	8dfe147710a63de848b9f48ef35df969636c70426bf3f432dc2249b2	Relocate
f6fa12bf-eecc-43d7-b4a6-961f07c36b20	8dfe147710a63de848b9f48ef35df969636c70426bf3f432dc2249b2	Relocate

Server Actions

Server	Action
air-6	hibernate
compute4	hibernate
compute3	hibernate

Figure 4-19: The Optimization plan generated by the Optimizer.

In detail, the Optimizer generated a plan containing three types of actions, namely:

- 3 VM Migrations from server air-6 to server air-5;
- 39 VM Relocations (in Figure 4-19 the VM Relocations have been truncated for reasons of space and picture clarity. A complete list is provided in Table 4-14);
- 3 server hibernations as a result of emptying the servers air-6, compute3 and compute4.

Based on the predictions as to the VMs CPU utilization in the next hour, the Optimizer predicted that when applying the plan, the energy consumption would be reduced by 60.9% whereas the (pricing-model dependant) revenue would be also reduced by 42.16%[4], outreaching the initial policy target set to a 37.5% of energy reduction. However, as discussed in 4.3.3.5, this was not an over-

[4] At this point an assumption was made that when a VM gets relocated, the revenue from operating the VM would be of the hosting DC, also to provide incentive to the other DC to host the extra load. Evidently, different pricing and/or synergetic business models would result in different revenue results.

provisioning of energy efficiency at the cost of reduced revenue; even if one server was activated, the policy target would not be reached.

Considering the above, this checkpoint was verified.

4.3.3.6.5 The Optimizer sends the generated plan to the Policy Maker to approve it

A first proof of the communication of the Optimizer with the Policy Maker is provided by Figure 4-19, where the optimization plan showcases an "Approved" status. Moreover, from the logs of the Policy Maker, one can easily detect the communication between the Optimizer and the Policy Maker:

```
[D 160911 11:14:43 optimizer:14] received optimizer request id: 64
[I 160911 11:14:43 web:1908] 200 POST /v1/policy/optimizer/approve (127.0.0.1) 1.49ms
```

The above snipped suggests that the Optimizer contacted the Policy Maker to approve the plan with id 64.

Considering the above, this checkpoint was verified.

4.3.3.6.6 The Optimizer forwards the plan to the Policy Actuator to implement the VM migrations and the server actions and to the Policy Maker to organize the VMs relocation

After the plan execution, we checked the logs of the eCOP Monitor DB in order to check for the outcome of the i) VM migrations and ii) server hibernations. The result of this inspection is presented in Figure 4-20[5].

Time	Component Type	Component ID	Container Type	Container ID	Action
2016/09/11, 14:58:15	server	compute3	rack	-	hibernate
2016/09/11, 14:58:14	server	compute4	rack	-	hibernate
2016/09/11, 14:58:13	server	air-6	rack	-	hibernate
2016/09/11, 14:58:05	vm	fbcf54e6-304f-415e-8a8c-54f6ca8d5707	server	air-5	migrate
2016/09/11, 14:57:58	vm	f692d5fd-e71e-4ba8-b930-5d382062a743	server	air-5	migrate
2016/09/11, 14:57:50	vm	eecb9cab-71c6-4eb0-921e-90d27c5b5c7a	server	air-5	migrate

Figure 4-20: The intra-DC optimization actions logged by the Policy Actuator to the eCOP Monitor DB.

As can be easily deduced, the VMs migrated were the ones included in the optimization plan (see Figure 4-19, as were the physical servers that were shut down. Moreover, it can be observed that the VM migrations began almost 45 minutes after the optimization plan was issued. This is due to the fact that the actions implemented by DOLFIN are successive, as follows:

[5] The logging dashboard of the eCOP DB presents timestamps in local time (at the time of writing this was EEST, GMT+3)

1. VM relocations;
2. VM migrations;
3. Server actions

This successive action implementation has been adopted to achieve maximum resilience in case an action fails to be successfully applied.

Next, contents of the DOLFIN Info DB as to VM relocation activity are presented in Table 4-14.

Table 4-14: DOLFIN Info DB VM Relocations log.

ID	SOURCE	SOURCE_UUID	TARGET	TARGET_UUID
5	RegionTwo	12d267ba-8a64-495d-9bbc-5a803f575a36	RegionOne	61f50e0d-2c22-464f-b8a8-931aff79e8f7
6	RegionTwo	14b8582b-d2fd-4a5b-80a3-cfde50056348	RegionOne	1839475b-c38c-4063-aac5-1bc070c4dd36
7	RegionTwo	2018a321-d141-4deb-993b-33b619f13bea	RegionOne	ddd1837e-620b-4031-9baf-1c69737dfa6a
8	RegionTwo	272a4fcb-61ab-4dc7-9983-009c31133884	RegionOne	0827b852-dbb6-428e-8b9b-cb926baa0d07
9	RegionTwo	27ee6efb-ad76-4c7a-b73f-06b97d44d67b	RegionOne	046516e9-8eab-4cfc-a768-c73a81992da9
10	RegionTwo	2b909a4b-9e39-4e24-ab4c-60077199eeb7	RegionOne	01a13197-e7a3-4314-805c-54d384e50148
11	RegionTwo	2caf0782-94ee-4519-ac41-0a9c2fdee152	RegionOne	161949f7-b405-4c0b-9aa1-e57706937d86
12	RegionTwo	31301098-526b-4e01-a2a3-8b9f2acdb678	RegionOne	2c00ca32-f36f-4445-bf8f-1cdbfcd5d7bd
13	RegionTwo	3272a1bb-eb68-498b-9325-e30c4ec57e3d	RegionOne	7314bee0-658c-440a-96af-f36c51294db7
14	RegionTwo	362d3dbe-a19d-4d9d-b179-db893c39dfa5	RegionOne	831fc649-4e99-4cb0-bbf5-9c1addd82133
15	RegionTwo	3d0c58e1-4ffc-4039-8da0-37e50b57f0a2	RegionOne	ba49fc36-fd97-43ca-9209-37639110a9a8
16	RegionTwo	3ebbe814-7c11-4f9a-8809-989bbe226655	RegionOne	33bb1302-60dc-4fd3-9ffe-27116f784656
17	RegionTwo	4328971c-8088-4c26-aed0-89dfe3a8fa1b	RegionOne	53b7f2a1-309c-4d4f-961e-09a199e6d53b
18	RegionTwo	4de30300-4c29-4bf7-8351-37dc9eaa1c47	RegionOne	a6d1ac5e-f603-4a9c-a8eb-42e3959f08a6
19	RegionTwo	59931547-204e-499b-9461-fc66de284a55	RegionOne	7ac4bb24-2747-4b52-95dc-3b1878a3f124
20	RegionTwo	5bdf20ef-7be0-411b-92d2-769cbd54c52c	RegionOne	30ee068e-7c85-45a7-a912-b208583f4c38
21	RegionTwo	5dc7d3cb-9bee-47e9-91b5-6fe0ea14d819	RegionOne	3f409cbf-2a30-4906-a5ec-084e6d76a2d1
22	RegionTwo	5dd87da1-03be-4038-b265-a88c9985b683	RegionOne	2f25281e-707e-4f71-9885-c70a4dbe7b5c

23	RegionTwo	63791238-6fe4-4fda-8d3d-031c84ed4301	RegionOne	ca979c1c-7a7c-4ccb-b08a-84fbdede69f0
24	RegionTwo	6462ee85-0383-44c2-abb6-8b26bc855520	RegionOne	cedaf029-da2c-4fab-a97d-3a51a994c5a2
25	RegionTwo	67cde55c-ad14-487f-8195-a4f4dd72e426	RegionOne	2478f71e-3e34-4851-8548-781550a77f39
26	RegionTwo	73e51164-fb6d-4b2f-992c-da79d9e9e8f8	RegionOne	a601d8c6-f5a3-41ba-8308-8c0046062a14
27	RegionTwo	758137d8-059e-4f9a-b193-3d3198f4f111	RegionOne	c6ece5a5-ce45-45af-8a95-f0435db402f8
28	RegionTwo	7edd4b56-48ee-476e-a458-229d9231613f	RegionOne	d2aea070-c95c-4d40-8137-2ef88d7ce9f6
29	RegionTwo	800f21cb-74e2-4b3e-8be8-fb851daff32c	RegionOne	060463fa-1f87-44dd-b182-fc884b2d064f
30	RegionTwo	8013aa0b-99a9-432a-992f-8573f3fb0ddd	RegionOne	56b5d78e-0a46-4634-a194-76b8c771e35a
31	RegionTwo	878e4d34-78b3-4aaa-b74b-f4dc64fd8b9a	RegionOne	fbdbfa13-fb84-4f30-9be5-d3f9b138f10a
32	RegionTwo	920b2be0-2f26-40fb-aa31-4916c70f209b	RegionOne	0181d410-aa2b-4ef6-a8ae-2df421981884
33	RegionTwo	94cb3e4e-173c-4571-8366-22340676754e	RegionOne	25acfac4-9e31-4bd5-8598-06b8161ef0e7
34	RegionTwo	abd40cc9-1131-41e5-8459-0dc5fc7dad1e	RegionOne	1157d8f7-aaea-487a-83da-75220d85390d
35	RegionTwo	b60b127b-4ada-4f28-9b88-f3e711f48327	RegionOne	fd02626b-2cb9-4447-b4d5-f3b23c4d8981
36	RegionTwo	b7dbec6c-2c40-4154-8a1f-ad21b8d11888	RegionOne	9e957bcc-4208-43d7-adab-6b84da6bc744
37	RegionTwo	b91e13cf-d812-4262-a3d5-7c39927032a0	RegionOne	f2c80b54-eb79-414f-9222-22df91bba918
38	RegionTwo	c37360ec-741d-445b-b1a0-16a4c77ef3fe	RegionOne	373f748d-3df0-49b7-b3f8-e523bf581284
39	RegionTwo	c77e98c8-7d2d-46b0-b85e-be6e116ee653	RegionOne	5450a3e6-d3b5-43ff-b426-b0ba59cf5320
40	RegionTwo	e2bf926a-65e6-4ecf-b72f-a6a725e19d24	RegionOne	f2769969-2c2d-4dc9-a974-ee2cf6408e8f
41	RegionTwo	e2eacc7f-27ab-4a48-8006-fea76a1ad442	RegionOne	d0baff69-ad8f-490b-bb91-067bd54a152a
42	RegionTwo	e6908446-987d-4054-a478-58ebd4d7a50e	RegionOne	6a79c088-a8c9-4752-9307-4f5da18c88a4
43	RegionTwo	f6fa12bf-eecc-43d7-b4a6-961f07c36b20	RegionOne	cc1b9ae2-4440-4999-91b8-9efe736da0c1

Note that when a VM gets relocated, it is assigned a new UUID hence, in order to be able to keep track of a VM across the various DCs, the UUIDs at both the source DC and the target one are presented. The timestamp of each relocation is also logged but was omitted in the above table for reasons of clarity and brevity.

Considering the above, this checkpoint was verified.

4.3.3.6.7 Check that the energy consumption of the DC has fallen at or under the requested limit.

Figure 4-16 provides evidence of the fact that the DC achieved its energy consumption goal. It is worth noting that there was no over-provisioning in the VMs relocations and server hibernations in the generated plan; in paragraph 4.3.2.5, evidence of why the energy consumption could not be higher than the achieved one (else the targets set would not be achieved) is presented.

4.3.3.7 Analysis of cost and energy savings

In addition to the core integration test 4.3.3, another accompanying test was conducted in a similar context, order to more deeply analyse benefits coming from the integration in a data centre of DOLFIN tools, and renewable energy sources (in particular confirm and validate the green power usage). The relevant tests took place in the PSNC testbed which features a set of photovoltaic panels used for collecting solar energy and providing it to the DC infrastructures. Table 4-15 presents the VM flavours defined for the green power analysis test and Table 4-16 contains information related to the resource allocation along the OpenStack Compute Nodes. The regions configuration is the same as mentioned in the section 4.3.3.3.

Flavour name	VCPUs	RAM	Disk[GB]
dolfin_mikro	1	1024	4
dolfin_small	2	1024	4
dolfin_medium	4	1024	4
dolfin_large	6	1024	4

Table 4-15: VM flavours of the green power testing scenario.

Compute	# of VMs	No. Cores	Ram[GB]
RegionTwo Compute 1	6	9	9
RegionTwo Compute 2	31	96	31
RegionTwo Compute 3	24	77	24
RegionTwo Compute 4	30	92	30
RegionOne Compute 1	0	0	0
RegionOne Compute 2	0	0	0
Total	91	274	94

Table 4-16: VM allocation before migration process

The test started before sun rise, before any green energy will be produced (total power production would be equal to 0). As time went by, the green energy production was expected to rise and the RegionOne server was expected to become “green”, namely powered by green energy sources (the

PVs). The total energy usage by RegionOne was expected to be, finally, supplied by the green energy source entirely. In detail, on test completion, the testing scenario was expected to demonstrate:

- The power production changes related to the different time periods within the day;
- The migration process triggered by energy production rise;
- Servers without any active suspension after the migration process;
- Cost reduction after migration load to the green server;

Figure 4-21: presents power generated by Photovoltaics system. The power drop observed at 8.50-8.55 was a result of rain clouds passing above the Photovoltaics panels. Figure 4-22: presents regions expenses related with running VMs. For the purposes of the test the cost of the 1kWh energy production was assumed to be equal 30 Euro cents. The cost of running RegionOne dropped down to 0 Euro cents per hour after the sun had risen; the rising power production triggered the migration process in the test bed. Figure 4-22 until Figure 4-26 present the changing load of the physical servers running compute nodes during the migration process.

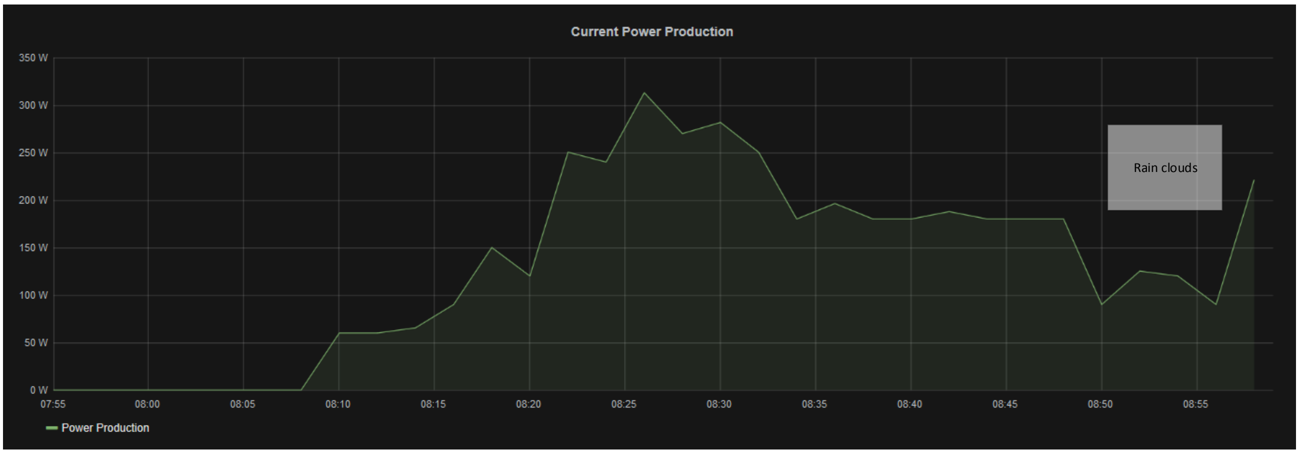

Figure 4-21: PV panels power production in the Morning during cloudy day

Figure 4-22: Comparison of power cost between RegionTwo and RegionOne ("Green" Power)

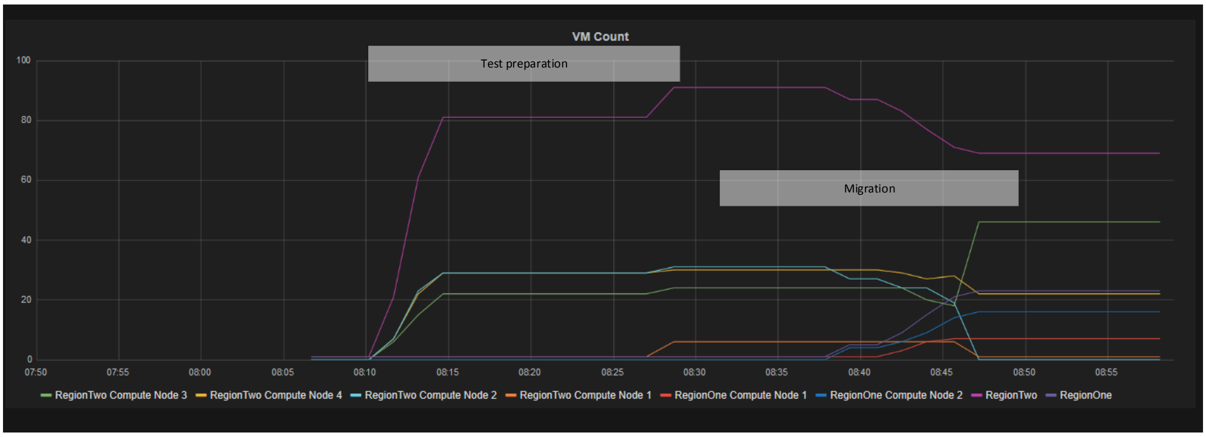

Figure 4-23: Number of VMs allocate in the Compute Nodes before and after the migration process

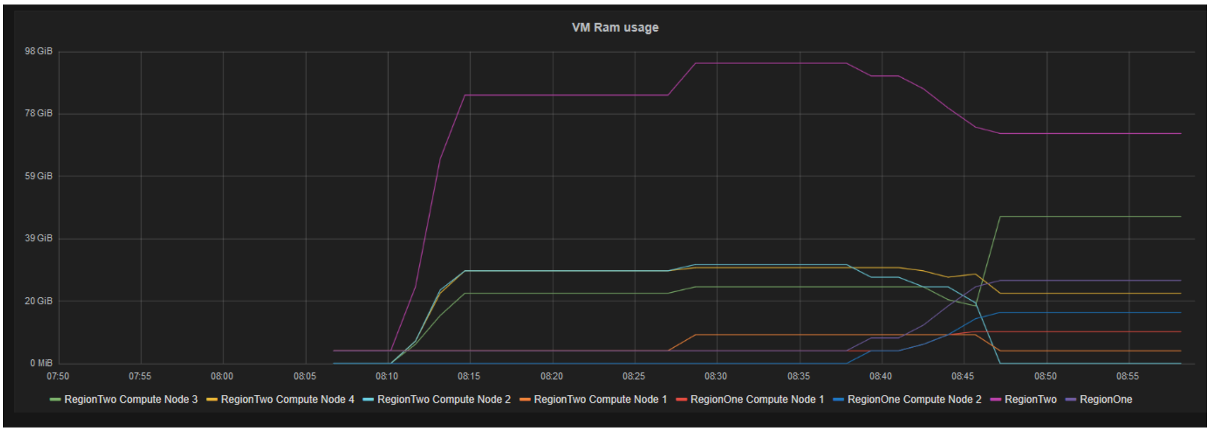

Figure 4-24: Total VM ram usage in the Compute Nodes before and after the migration process

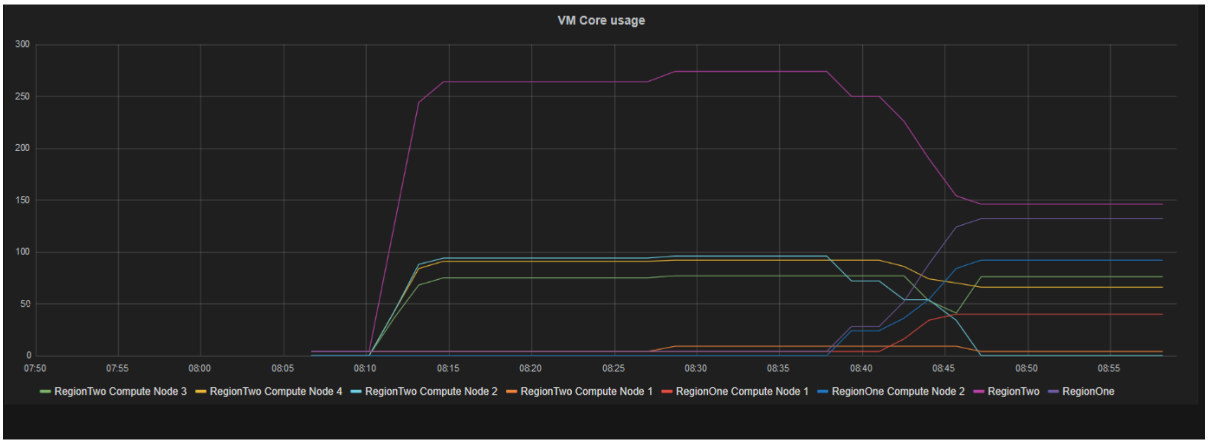

Figure 4-25: Total number of cores assigned to VMs in the Compute Nodes before and after the migration process

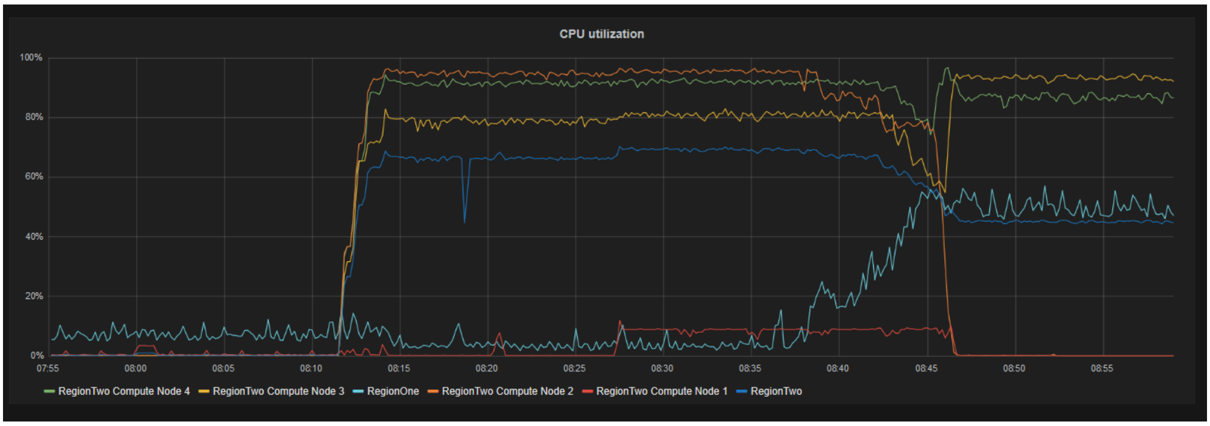

Figure 4-26: CPU utilization on the particular nodes before and after the migration process

Table 4-17 presents VM allocation after finalizing the migration process. After the migration process both Compute node 1 and Compute node 2 were suspended. The green energy production was sufficient for powering the RegionOne. The cost of the RegionOne maintenance decreased about 42%. CPU utilization in RegionOne dropped by 32%.

Table 4-17: Final VM allocation

Compute	# of VMs	vCPUs	RAM [GB]
RegionTwo Compute 1	1 (suspended)	4	4
RegionTwo Compute 2	0	0	0
RegionTwo Compute 3	46	76	46
RegionTwo Compute 4	22	66	22
Summarized RegionOne	23	132	26

The above testing scenario indicates another key aspect of DOLFIN capability to exploit the different energy state of various DOLFIN-enabled synergetic DCs; when a DC being powered by own sources (RegionOne) is able to host IT load from other DCs (RegionTwo) and accordingly changes the active DC Policy to allowing cross-DC workload, significant overall brown energy consumption benefits can occur in practical terms in aggregate simultaneously ameliorating the energy mix and allowing for the enforcement of green SLAs.

Figure 4-27 presents a comparison of power production between different weather conditions, including:

- cloudy day with chassis average load (chassis power usage around 1.0-1.1kW);
- sunny day with chassis average load (chassis power usage around 1.0-1.1kW);
- sunny day with chassis high load (chassis power usage around 2-3kW).

All energy produced by the PV system was consumed immediately by chassis. Only during cloudy day in the Morning the power production was not sufficient to supply power requirements of the chassis, for the rest measurements the power production was equal to power consumption.

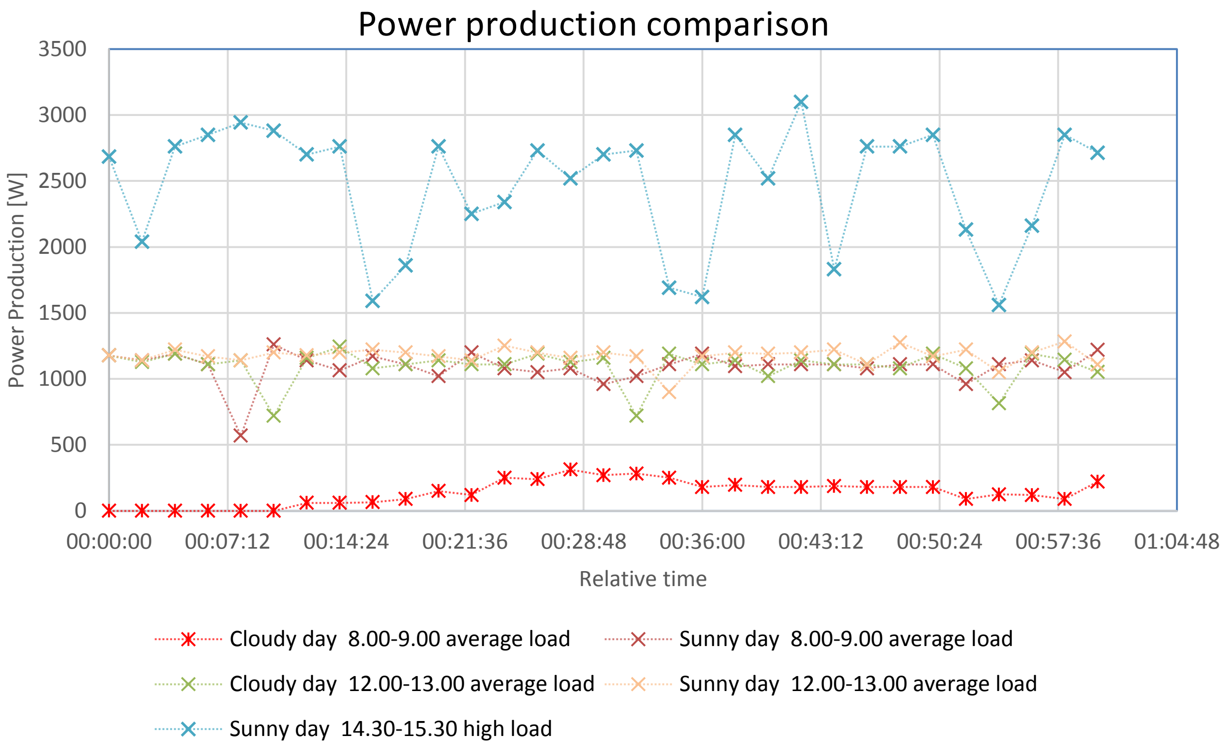

Figure 4-27: Comparison of power production between different weather conditions.

5 Evaluation of DOLFIN as a whole

5.1 Emulated Evaluation

In this section, the evaluation of the DOLFIN platform as a whole is presented on the basis of the DC flavours presented in Section 2. The evaluation process was implemented as an emulation in order to be able to assess the expected performance of DOLFIN on the largest possible scenarios and DC configurations, without being limited on the trial DCs tested under the DOLFIN project scope.

5.2 Evaluation Procedure

The evaluation procedure can be summarized as a set of semi-random DC configurations containing a (each time) variable numbers of DC rooms, racks, servers and VMs, the number of the latter changing as a function of average server CPU and/or RAM utilization. In this sense and for the rest of this section, the term "DC configuration" will refer to a static number of:

1. DC rooms;
2. Racks per DC room;
3. Servers per rack;
4. CPU / RAM utilization[6].

For each configuration, twenty (20) emulations were performed in order to reduce the effect of randomness in the generation of VMs and servers' characteristics and the respective measurements. In each emulation, the characteristics (number of logical CPU cores, memory capacity, HDD capacity, consumption characteristics and whether it is a green one or not) of the servers change randomly bound by preconfigured minimum and maximum values, altering the DC configuration as to its computing power and energy consumption characteristics. Similarly, each time, four (4) different VM flavours (e.g. VM virtual hardware configurations) are defined in a random, partially preconfigured, manner and the instantiated VMs follow the specifications of one of the generated flavours, at

[6] For the present evaluation only average CPU utilization has been considered, though setting RAM utilization is also allowable through the evaluation framework settings.

random. Based on the VMs characteristics, a set of semi-random[7] measurements are being determined. In order to enable the rest of the DOLFIN components to be able to operate properly on the basis of these data sets, all this information is being stored in the eCOP DB.

Considering the above, the evaluation framework performs the following operations in each emulation:

1. Sets the policy for the optimization, operating in the form of the Policy Maker. The two policies considered for the present evaluation set contain energy-efficiency- and performance-oriented policies, allowing cross-dc workload (see D3.4 [26] for details and discussion).

    a. The policy tuned for energy efficiency primarily cares for minimizing the overall DC energy consumption;

    b. The policy tuned for performance attempts to make better use of the most high-performing IT infrastructures in order to provide services of better quality, hence revenue based on performance-based SLAs (see section 3 for details on the revenue model assumed for the purposes of the evaluation).

2. Clears the state of the eCOP DB by deleting all relevant DC entities and measurements;

3. As a next step, it creates the VM Flavours that will be supported throughout the emulations.

4. Then, it creates the DC rooms, starting from the MIN_ROOMS_NO configuration option.

5. Next, the racks for each room are generated, starting from the RACKS_MIN_PER_ROOM_NO configuration option which sets the number of racks per DC room. Considering the previous step, initially MIN_ROOMS_NO × RACKS_MIN_PER_ROOM_NO racks will be created.

6. Next, the DC servers get generated, starting from the SERVERS_MIN_PER_RACK configuration option, which sets the number of servers per DC rack. In the first loop, MIN_ROOMS_NO × RACKS_MIN_PER_ROOM_NO × SERVERS_MIN_PER_RACK servers will be created.

7. For each server, the script sets the initial maximum workload based on the MAX_INITIAL_UTILIZATION_SERVERS_CPU configuration option.

8. Next, the VMs get generated as follows: a random flavour gets selected and a VM instance is considered based on this flavour. For the current server, the total load (in terms of CPU utilization) got calculated and if it did not exceed the maximum allowed workload, it got assigned to this server, else it got dropped. It should be highlighted that, based on the random flavour selection, the requirements for the VMs were not homogeneous; therefore, the number of VMs per server (thus in total), was not standard per emulation.

[7] The semi-randomness is based on the following: for each VM, a pseudo-random numerical ID gets generated and is fed to a sine function to affect the respective period. Next, based on the current emulation time, a sine value between 0 and 1 is calculated and is multiplied by the CPU/RAM characteristics of the VM, as dictated by its flavor to get the semi-random, to get the CPU/RAM measurements.

9. A number of OPTIMIZER_REPETITIONS_PER_SETUP emulations for each setup gets carried out.

10. After the execution of the OPTIMIZER_REPETITIONS_PER_SETUP emulations has been completed, the script increases the server utilization profile, number of servers, number of racks and number of rooms to achieve a complete emulation. In this sense, a total of

$$\begin{aligned}\text{EMULATIONS\_NO\_TOTAL} &= (\text{MAX\_ROOMS\_NO} - \text{MIN\_ROOMS\_NO}) \\ &\times (\text{RACKS\_MAX\_PER\_ROOM\_NO} - \text{RACKS\_MIN\_PER\_ROOM\_NO}) \\ &\times (\text{SERVERS\_MAX\_PER\_RACK} - \text{SERVERS\_MIN\_PER\_RACK}) \\ &\times ((\text{MAX\_INITIAL\_UTILIZATION\_SERVERS\_RAM} \\ &- \text{MIN\_INITIAL\_UTILIZATION\_SERVERS\_RAM}) / 10) \\ &\times \text{OPTIMIZER\_REPETITIONS\_PER\_SETUP} \times \text{POLICIES\_NO}\end{aligned} \quad (1)$$

emulations will be performed.

Indicatively, for the emulation processes in the context of urban micro-DCs operating 2 – 5 racks, the core configuration options are tabulated and presented in Table 5-1, below:

Table 5-1: Basic configuration of the core emulations set for the case of urban micro-DCs.

Variable	Description	Value
MIN_ROOMS_NO	The minimum number of rooms per scenario	1
MAX_ROOMS_NO	The maximum number of rooms per scenario	2
RACKS_MIN_PER_ROOM_NO	The minimum number of racks per room	2
RACKS_MAX_PER_ROOM_NO	The maximum number of racks per room	5
SERVERS_MIN_PER_RACK	The minimum number of servers per rack	2
SERVERS_MAX_PER_RACK	The maximum number of servers per rack	6
SERVERS_MIN_RAM_GB	The minimum possible amount of RAM of a server in GB	16
SERVERS_MAX_RAM_GB	The maximum possible amount of RAM of a server in GB	128
SERVERS_MIN_CPU_CORES	The minimum possible number of logical cores of a server	16
SERVERS_MAX_CPU_CORES	The maximum possible number of logical cores of a server	24
SERVERS_MIN_FREQ	The minimum possible maximum frequency of a server in GHz	1.8
SERVERS_MAX_FREQ	The maximum possible maximum frequency	4.0

	of a server in GHz	
MIN_INITIAL_UTILIZATION_SERVERS_CPU	The minimum initial CPU utilization of the servers in total (DC utilization)	20%
MAX_INITIAL_UTILIZATION_SERVERS_CPU	The maximum initial CPU utilization of the servers in total (DC utilization)	80%
OPTIMIZER_REPETITIONS_PER_SETUP	The number of repetitions to optimize a certain setup (number of rooms, racks and servers under fixed initial aggregate DC utilization)	10
SERVERS_PERCENTAGE_GREEN	The percentage of green-powered servers	10%

Given the above and based on (1), for the above setup a number of about 7,000 emulations and (accordingly) optimization plans were conducted. The source code of the evaluation framework can be accessed online through the DOLFIN source code management platform [31], where details about the entire set of emulation parameters as well as instructions on how to configure the evaluation procedure are given.

The following figures offer an overview of a snapshot DC configuration as was valid for emulation number #912 of the evaluation process. The relevant configuration contained 1 DC room, holding 4 Racks, with 5 physical servers that hosted a total of 62 VMs. The mean CPU utilization of the physical servers was set to 20%.

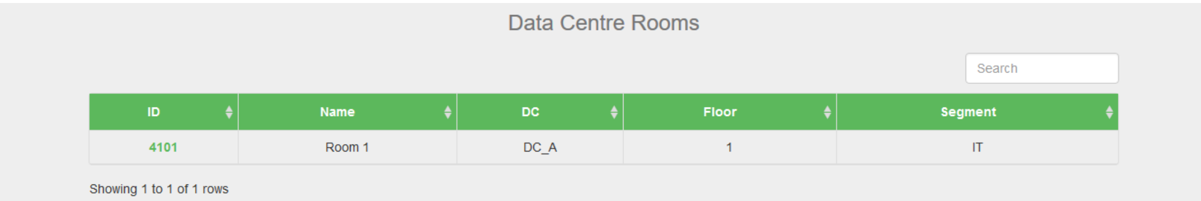

Figure 5-1: Snapshot of a DC Room considered in a single emulation of the evaluation (#912).

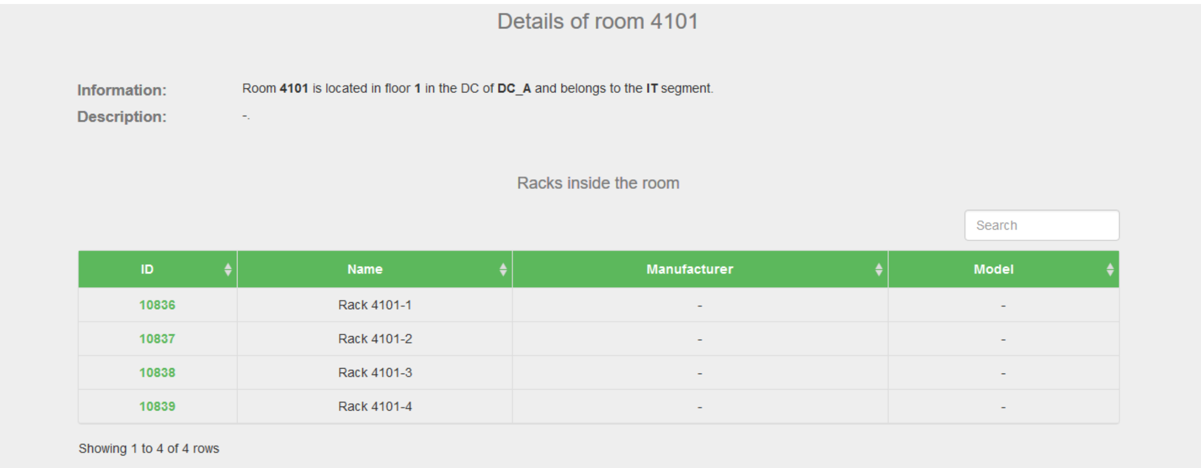

Figure 5-2: Snapshot of the DC Racks supported by the emulation #912.

Details of Rack 10836

Room: Rack Rack 4101-1 (10836) is located in room 4101.
Model: - -.
Description: -.

Physical servers currently hosted

Search

Serial Number	Name	Model	Active	Green
18931584-e9da-4ca6-a662-fd08b00536d3	Server 10836-4	- - -	True	False
47f4ea08-e1dc-454b-9c8d-41d3f2e984d7	Server 10836-3	- - -	True	False
5084f1a9-574c-437d-9115-f70cd3dfaec4	Server 10836-2	- - -	True	False
607d25a2-ad9d-4416-97b4-dd9f634d9af7	Server 10836-1	- - -	True	False
bbd09932-618b-4e45-85ff-6895d0b1b8b8	Server 10836-6	- - -	True	False
dc3874da-607b-4449-afaf-76f7fdffeb6c	Server 10836-5	- - -	True	False

Showing 1 to 6 of 6 rows

Figure 5-3: Snapshot of the servers of the DC Rack #10836 in the context of the emulation #912.

Details of Server Server 10836-4 (18931584-e9da-4ca6-a662-fd08b00536d3)

Info: - -, serial number is **18931584-e9da-4ca6-a662-fd08b00536d3**.
Description: -. The server is currently **active**.
Rack: The server is located in rack 10836.
Green: This server is not green.

Detailed Characteristics

Cores	Max. Frequency	RAM (GB)	HDD (GB)	Min/Max W per Core	Min/Max W per RAM MB	Min/Max W per kbps
16	3003 MHz	64	16384	21.00 / 45.00	- / -	0.00 / 0.00

Virtual machines currently hosted

Search

UUID	Name	User	Flavour	Status	Task	Instantiation Time
14f25be9-72f5-4755-8fb8-c8fb53291412	VM-7-Server 10836-4	-	3 CPUs, 5120 MB RAM, 23 GB HDD	ACTIVE	None	2016/09/08, 06:35:34
27747cac-bdec-4da0-b623-5983cbaa1501	VM-2-Server 10836-4	-	3 CPUs, 4608 MB RAM, 21 GB HDD	ACTIVE	None	2016/09/08, 06:35:33
8c4993d5-068b-4f8e-b4a9-7ef1b2b97028	VM-4-Server 10836-4	-	3 CPUs, 4608 MB RAM, 21 GB HDD	ACTIVE	None	2016/09/08, 06:35:34
93afdf50-d1fe-49dc-92d7-622eb1d9cb43	VM-8-Server 10836-4	-	2 CPUs, 6144 MB RAM, 35 GB HDD	ACTIVE	None	2016/09/08, 06:35:34
96356b35-5e17-487c-a0d8-6d86ecd9967f	VM-5-Server 10836-4	-	3 CPUs, 4608 MB RAM, 21 GB HDD	ACTIVE	None	2016/09/08, 06:35:34
a936ab11-cc0a-43ce-9786-303098901ecc	VM-9-Server 10836-4	-	3 CPUs, 4608 MB RAM, 21 GB HDD	ACTIVE	None	2016/09/08, 06:35:34
d6f8a7ab-cf7d-453f-978e-c71914d9dbd9	VM-1-Server 10836-4	-	3 CPUs, 6656 MB RAM, 33 GB HDD	ACTIVE	None	2016/09/08, 06:35:33
d79f8312-00b7-4794-86f0-643a648312e8	VM-6-Server 10836-4	-	3 CPUs, 5120 MB RAM, 23 GB HDD	ACTIVE	None	2016/09/08, 06:35:34
e830b853-0915-4813-b714-1bb16f63225e	VM-3-Server 10836-4	-	3 CPUs, 5120 MB RAM, 23 GB HDD	ACTIVE	None	2016/09/08, 06:35:34

Showing 1 to 9 of 9 rows

Figure 5-4: Snapshot of a server's characteristics and of the VMs hosted by it in the context of emulation #912.

5.3 Evaluation outcome

In this section, the core conclusions from the DOLFIN evaluations as a whole are drawn, based on the configuration tabulated in Table 5-1.

5.3.1 Urban micro-DCs

In this case, we examine micro-DCs containing a very limited number of racks (2 – 5) in a single room. In the following, the performance of DOLFIN is examined on the basis of the above configuration (see Table 5-1), generating a number of 11,200 different scenarios; 20 repetitions per DC setup were conducted. Although we conducted the evaluation with the number of servers per rack changing from 2 up to 6, next the results for 5 servers per rack are presented, to facilitate the comparison with the forthcoming medium-sized DC results.

The following figures present the average number of VMs, the energy gain and the revenue benefit from applying DOLFIN when the initial micro-DC load changes from 20% up to 80% and the optimization policy has been setup (by the Policy Maker) to optimize against the energy consumption.

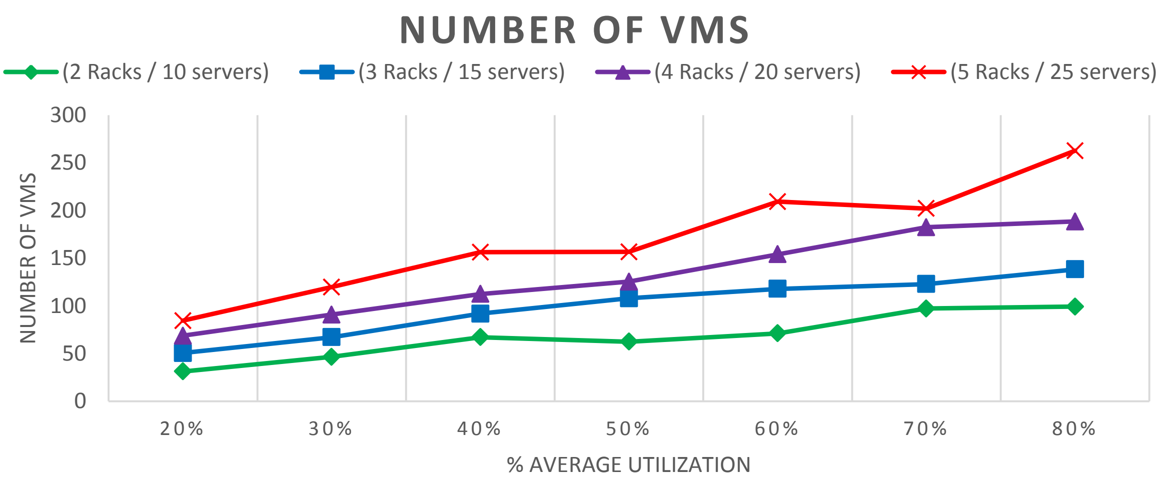

Figure 5-5: Average number of VMs hosted by the Micro-DC as a function of the average micro-DC utilization.

The results indicate that as the number of hosted VMs increases with the average DC utilization (Figure 5-5), the energy gain decreases as the DC configuration options get less and the possibility to result in inactive servers to put them in sleep state decreases (Figure 5-6). Interestingly, a simple linear regression analysis indicates that the rate of energy benefit reduction as a function of the average DC utilization is on average about 4.1% for each 10% of increase in the average DC utilization with an R^2 value of 0.96. In absolute numbers, the percentage of the expected energy gain for highly under-utilized micro-DCs was ranged between 65% up to 75% whereas the respective numbers for highly-utilized micro-DCs was much lower ranging from approximately 42% down to 37%.

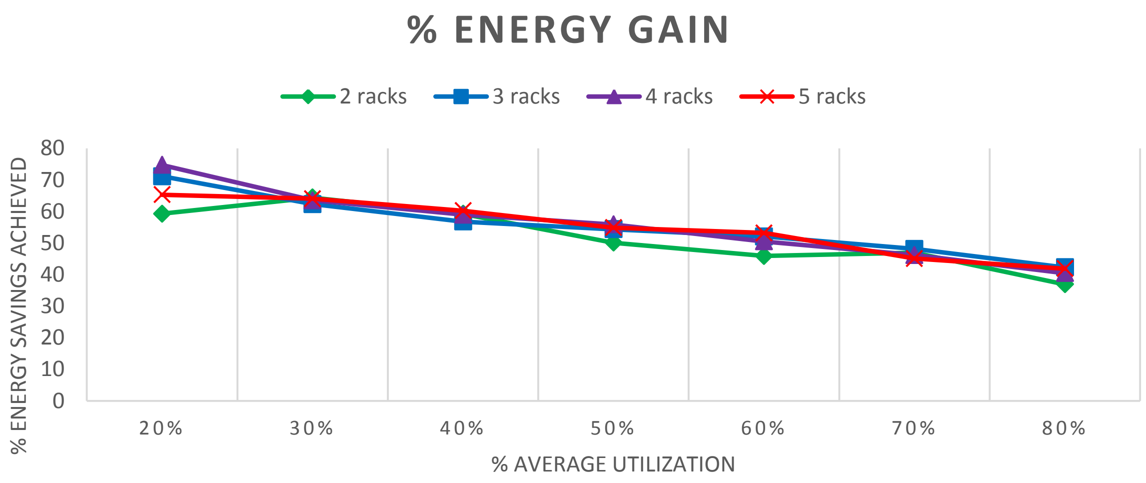

Figure 5-6: Average energy consumption benefit as a function of the micro-DC utilization (energy efficiency policy)

Simultaneously with the decrease in the energy benefit as the DC utilization increases, the expected revenue change also decreases (Figure 5-7), at a rate similar to the energy reduction one, presenting on average a benefit reduction rate of approximately 5.6% with an average R^2 value of approximately 0.9. In absolute numbers, the expected energy gain ranged from approximately 45%-55% for the case of highly under-utilized DCs down to 16%-18% for the case of highly-utilized DCs. This behaviour is to be expected as the limited number of servers (thus limited heterogeneity on the hardware energy consumption) in combination with the policy applied the optimization procedure which was set to reduce the energy consumption of the micro-DC, can explain the strong dependence between the two values.

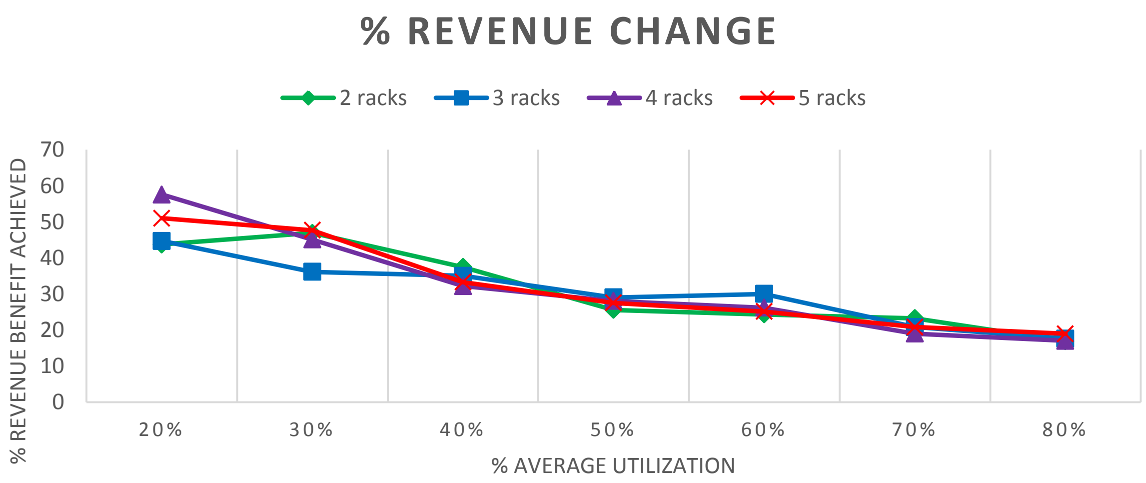

Figure 5-7: Average revenue benefit as a function of the average micro-DC utilization (energy efficiency policy).

Indeed, the following table presents the results of a simple correlation analysis between the expected energy and the expected revenue gain after the application of DOLFIN, the policy being set to optimize against energy consumption minimization. The high correlation value (0.91) indicates that the two attributes are highly correlated, the revenue benefit being caused by the lowered energy consumption of the DC elements.

Table 5-2: Correlation Analysis between the energy and the revenue gain for the case of 2 racks (energy policy).

	Energy Gain	*Revenue Gain*
Energy Gain	1	
Revenue Gain	0,91	1

Next and as regards the DOLFIN system operation under a policy set to maximize the performance (hence revenue based on the adopted revenue model detailed in Section 3). The following figures summarize the relevant findings.

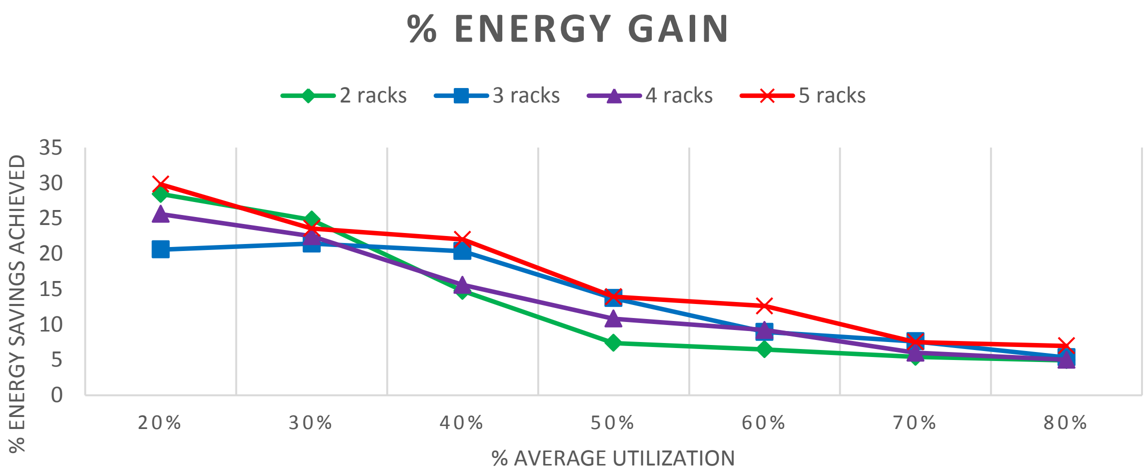

Figure 5-8: Average energy consumption benefit as a function of the average micro-DC utilization (performance policy).

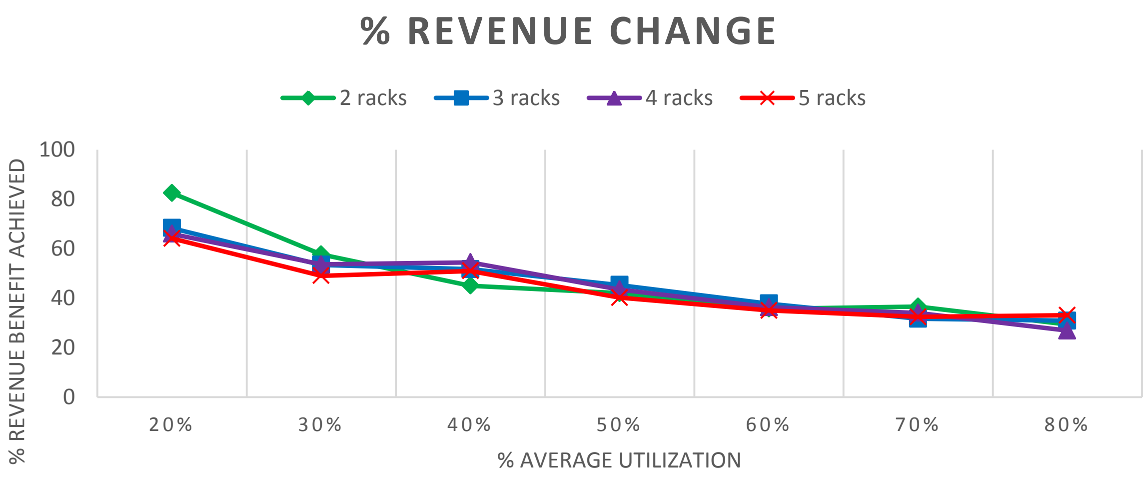

Figure 5-9: Average revenue benefit as a function of the average micro-DC utilization (performance policy).

As expected, as DOLFIN did not account for prioritizing energy efficiency during this set of emulations, the percentage of the energy gain attained was significantly lower than when optimizing against energy efficiency. Moreover, the rate of change as a function of the average DC utilization dropped to an average of approximately 3.6% for each 10% increase in the average DC utilization, indicating that when operating under the "Performance" policy, the average utilization is of lesser importance; it is actually the heterogeneity of the energy characteristics of the more powerful (from a computational point of view) IT equipment that plays the most important role. Simultaneously, the absolute numbers of the expected energy benefit also dropped in by 40% (in absolute numbers) compared to the energy efficiency-oriented optimization case. The following figure graphically presents the above.

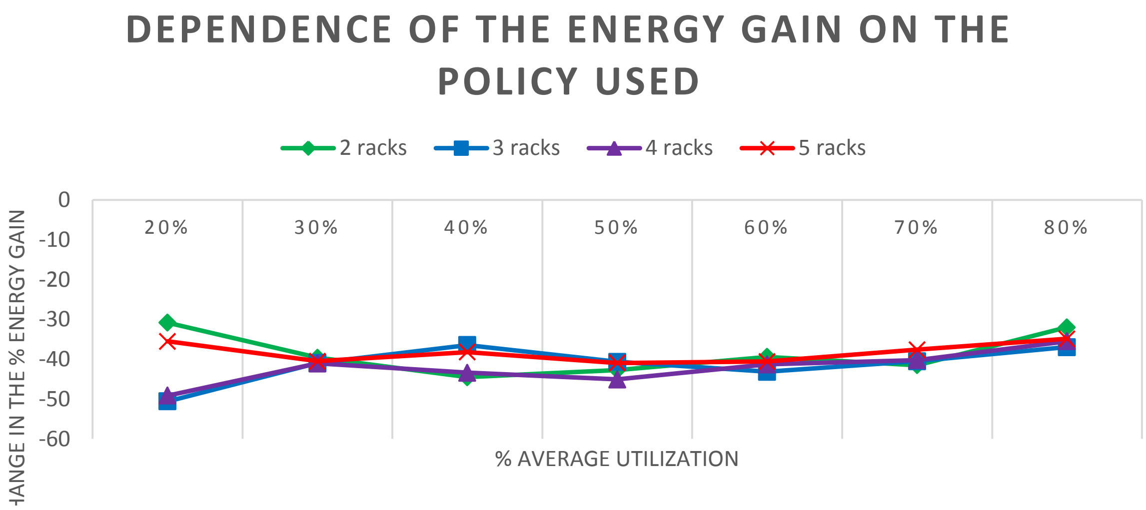

Figure 5-10: Difference between the expected energy consumption change when optimizing against performance instead of energy efficiency.

As apparent from Figure 5-10, the energy consumption reduction when optimizing against energy efficiency targets significantly outpaces the respective when optimizing for DC performance; the average change on the anticipated energy benefit is approximately 40.1%.

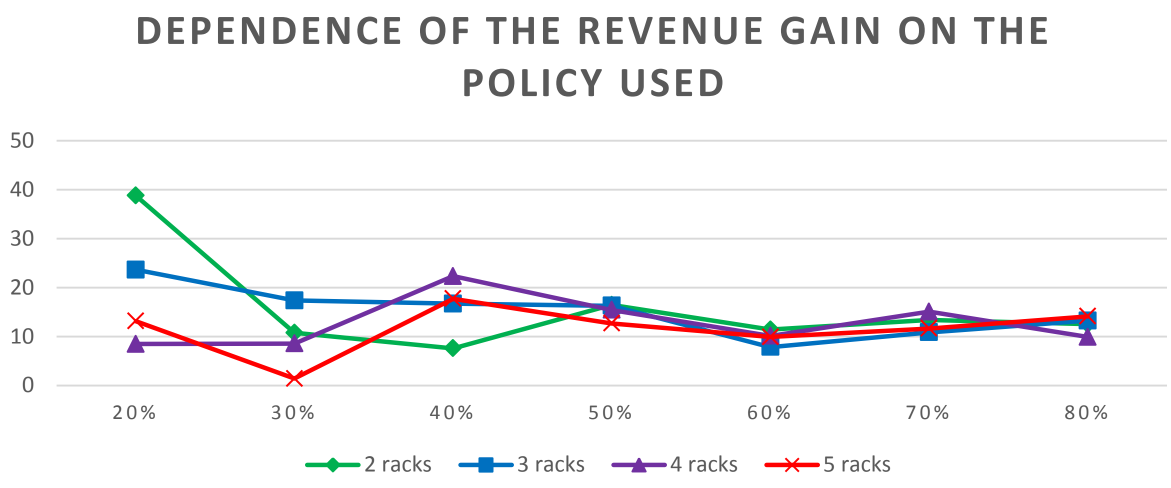

Figure 5-11: Difference between the expected revenue change when optimizing against performance instead of energy efficiency.

Despite the significant drop in the expected energy benefits, the expected revenue benefit in the two cases is on the positive side. This is presented in Figure 5-11, where the dependence of the revenue change on the policy used is depicted; despite the changes depending on the DC configuration and utilization, the average change on the anticipated revenue is, in total, 13.8%.

Finally, despite the change of in the absolute numbers of energy efficiency and revenue change when employing different policies, the correlation between the change in the energy gains and the expected revenue gains remains significant as deduced from Table 5-3.

Table 5-3: Correlation Analysis between the energy and the revenue gain for the case of 2 racks (performance policy).

	Energy Gain	*Revenue Gain*
Energy Gain	1	
Revenue Gain	0,87	1

5.3.2 Urban medium-sized urban DCs

In contrast to small-sized DCs that normally feature 2 – 10 racks with an average of 5 being the rule of thumb, medium-sized urban DCs are generally considered to contain 6 – 80 racks with an average number being 25. In the following, the performance of DOLFIN is examined on the same basis as the above configuration, though the number of racks was configured to be between 20 and 50, generating a number of 860 extra simulations; as in the case of the microDCs, 20 repetitions per DC setup were conducted. The number of servers per rack was fixed to five (5), hence the cases of 100, 150, 200 and 250 servers were considered. Last, for this set of emulations we considered that approximately **10% of these servers** (chosen at random) were **powered by green sources**, thus not contributing to the overall DC (brow) energy consumption.

The following figure presents the number of VMs that were considered, on average, for each DC setup:

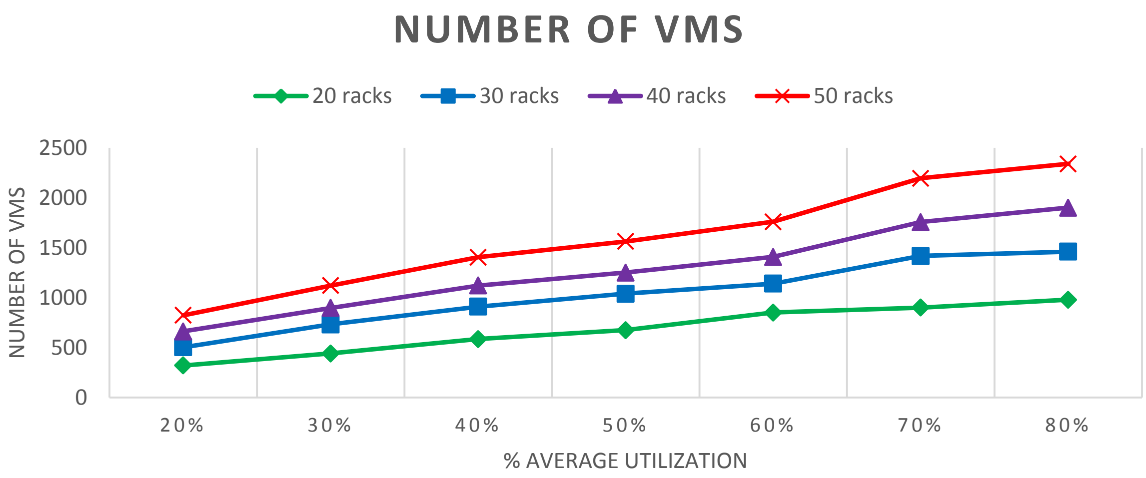

Figure 5-12: Average number of VMs hosted by the Micro-DC as a function of the average medium—sized DC utilization

Next, the we set the DC policy to optimize against energy efficiency and calculated the average energy and revenue benefits that were expected to be acquired after the application of the generated optimization plans. The results from this set of emulations are presented in Figure 5-13 and Figure 5-14.

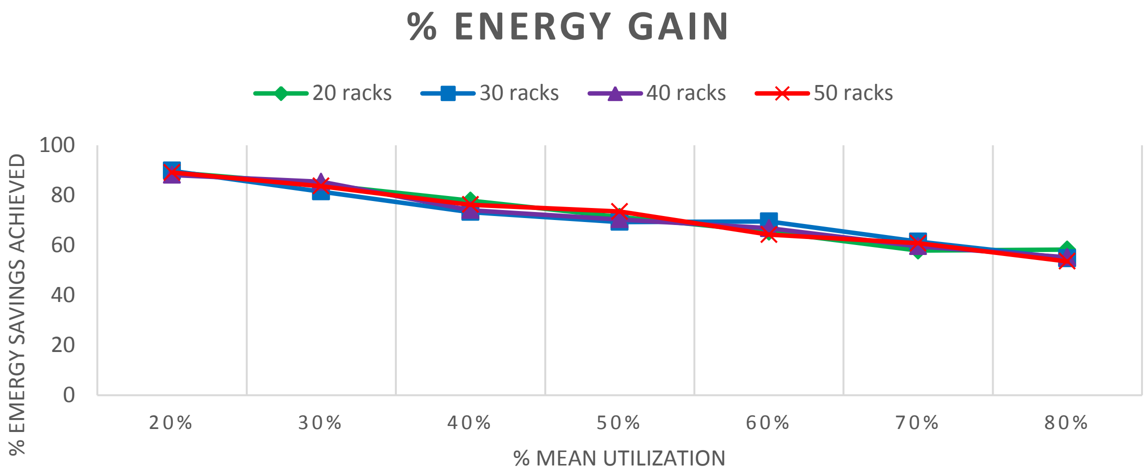

Figure 5-13: Average energy consumption benefit as a function of the average medium- sized DC utilization (energy efficiency policy).

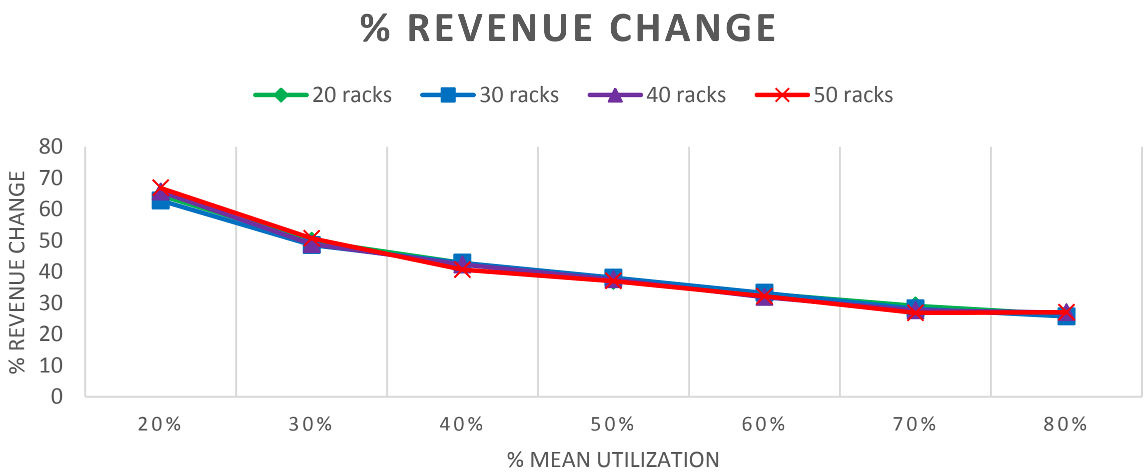

Figure 5-14: Average revenue benefit as a function of the average medium- sized DC utilization (energy efficiency policy).

Comparing with the micro-DC case, it can be easily deduced that average expected energy benefit as seen in Figure 5-13 is significantly larger than in the case of the micro-DCs when the DC is highly-utilized; this can be attributed to two things:

1. The increased number of available configurations, as it allows for more proper optimization, exploiting at the highest possible extent the heterogeneity of the servers available;
2. The introduction of the green-powered servers, as the Optimizer attempts to load them as much as possible;

Interestingly and in accordance with the aforementioned 2 differentiating factors, the rate of energy benefit reduction (as calculated through a simple linear regression) is approximately 5.3% per 10% of increase in the mean DC utilization (with an R^2 value of 0.96), which is significantly lower than in the case of micro-DCs indicating that DOLFIN can yield very significant energy benefits in partially green-powered medium-sized DCs even if they are highly-utilized.

Similar remarks hold for the average expected revenue benefit from the application of DOLFIN, where the expected revenue reduction per 10% of increase in the average DC utilization is approximately 5.77% with an R^2 value of 0.94 (on average).

As regards the correlation between the energy efficiency benefit and the revenue benefit, a very high correlation indicator was calculated reaching 0.97. Next, we set the DC policy to optimize performance and the respective results are drawn in the figures following.

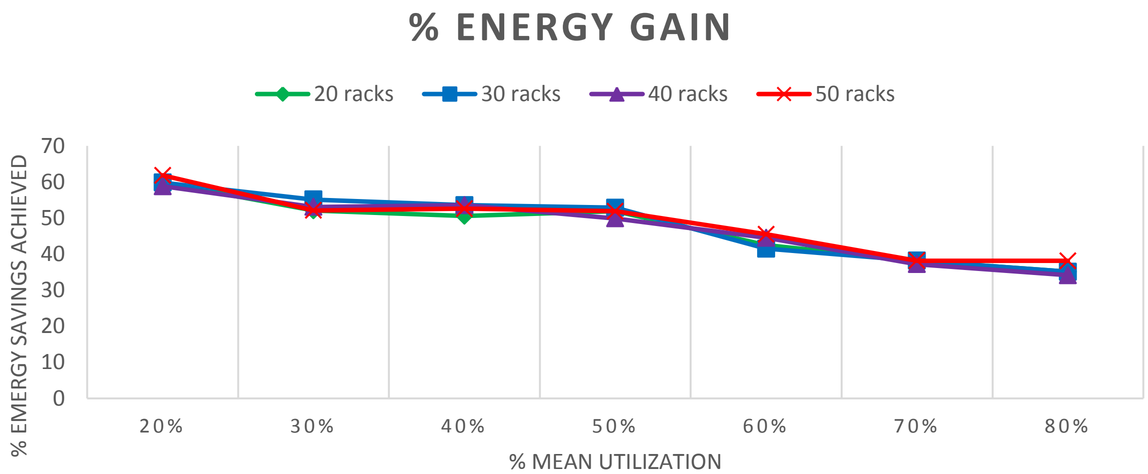

Figure 5-15: Average energy consumption benefit as a function of the average medium- sized DC utilization (performance policy).

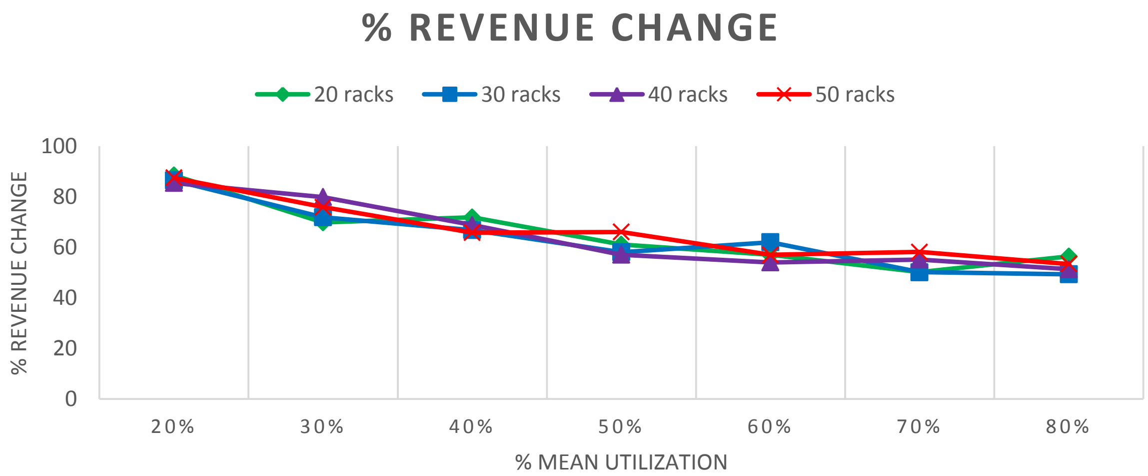

Figure 5-16: Average revenue benefit as a function of the average medium- sized DC utilization (performance policy).

As can be easily deduced by a simple examination of Figure 5-15 and Figure 5-16, once again, the actual scale of the DC in terms of number of racks does not significantly alter the performance of the DOLFLIN optimization. The average rate of change of the energy benefit with respect to 10% of increase in the mean DC utilization is on average approximately 3.8% with an R^2 value of 0.90. Similar deductions can be made also for the examination of the expected revenue change as the average DC utilization increases, the respective rate of change being on average 5.24% with an R^2 value of 0.89.

Based on the above results, Figure 5-17 depicts how the expected energy benefit changes as the policy changes (from energy efficiency-oriented to performance-oriented), whereas Figure 5-18 presents the respective results focusing on the expected revenue change.

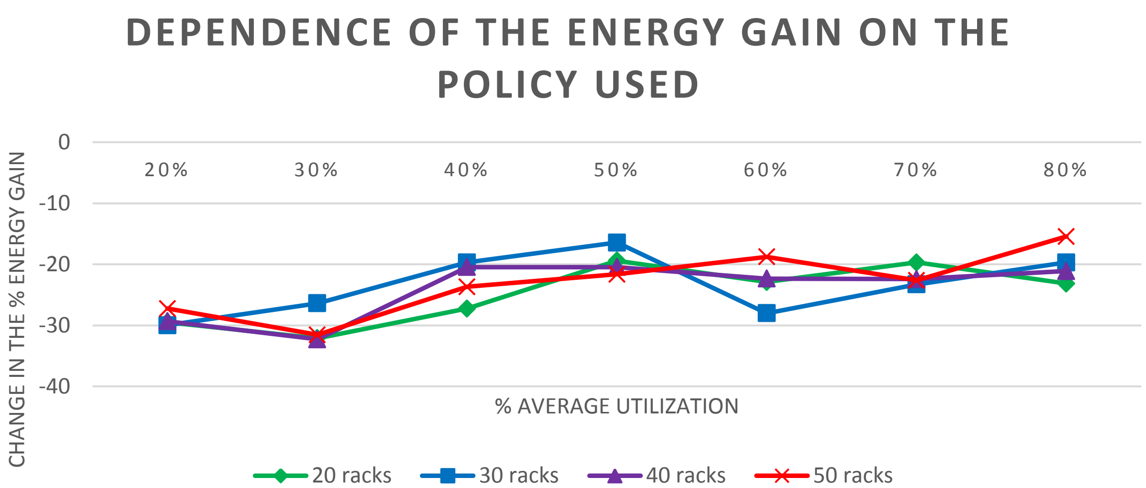

Figure 5-17: Difference between the expected energy consumption change when optimizing against performance instead of energy efficiency (medium-sized DC).

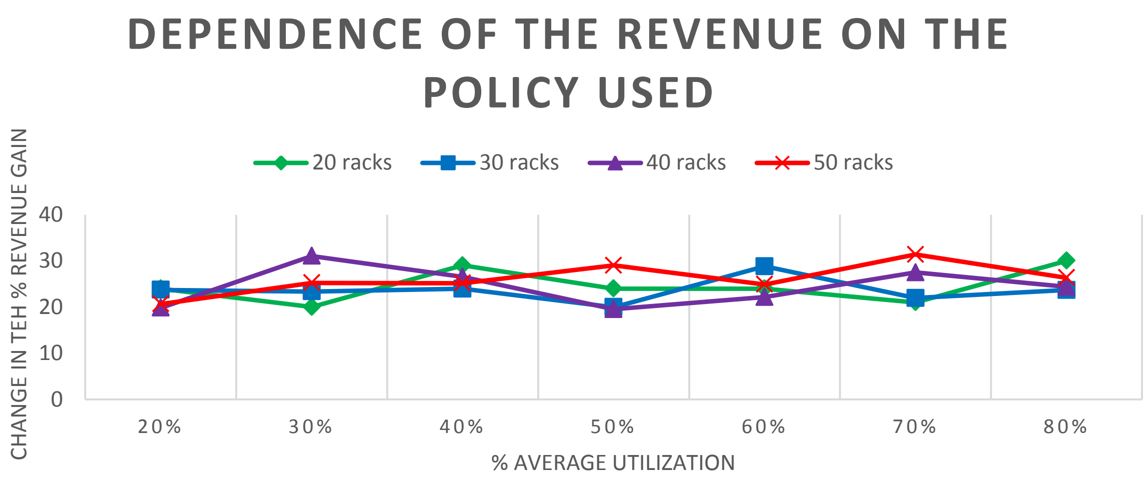

Figure 5-18: Difference between the expected revenue change when optimizing against performance instead of energy efficiency (medium-sized DC).

As expected, the expected energy benefit was significantly lower for the case of performance-based optimization (hence the negative values in Figure 5-17) while, simultaneously, the expected revenue was significantly increased as a result of the higher value of the services offered (increasing the price of the service provided based on the revenue model assumed and described in Section 3). The average expected drop in the expected energy revenue reached -23.8% whereas the respective increase for the revenue metric was 24.68%.

Comparing the medium-sized with the micro-DC case, it is evident that the DOLFIN impact on the latter is much more significant than in the former. This can be attributed to the increased flexibility that the increased number of physical servers delivers. Of course, the presence of the green-powered servers also plays a very important role in the exhibited increase of the expected energy and revenue benefits identified in the context of the medium-sized urban DCs. However, it should be highlighted that despite the context, it is easily deduced that:

1. The DOLFIN solution is able to scale to both types of urban DCs;
2. Its performance is almost linear to the level of average DC utilization;
3. It is able to be configured to favour energy efficiency over performance (hence revenue) and vice versa.

It should be also noted that;

1. A different pricing/revenue model would result in different evaluation outcomes in terms of revenue analysis;
2. In actual smart grid conditions where the price changes due to the Smart Grid Operator instructions might be higher, different results would be also attained.

5.4 Scalability and Stability of the Information Flows

This section provides a scalability evaluation of the DOLFIN S/W Information Service (IS) running on the VLSP testbed at UCL. First, we detail our experimental setup, relevant methodological issues, the performance metrics we used, plus our experimental scenarios. Then we present the experimental results from these scenarios, showing data from runs with 500 virtual nodes.

Each experimental run started with creation of a new virtual network topology being deployed at all physical servers. The topology consists of the number of Virtual Nodes, specified in each run configuration, and a number of virtual links being created randomly. The link details are picked from a distribution (i.e. a discrete distribution with a minimum of one, to maintain connectivity).

For each experiment we have created our own management probes associated to most appropriate node with diverse requirements in terms of information handling, including applications collecting information from the virtual nodes and applications requesting information from the DOLFIN platform. The management probes periodically transmit performance measurements to the Information Service over the management information flows. We performed tests with management probes deployed at the virtual nodes or as standalone physical applications. The former was used for controlling the virtual network and the latter for controlling the physical infrastructure.

We carried out experiments highlighting aspects such as the flexibility, and the stability / scalability behaviour of the Information Service based on the following scenario:

Scalability / Stability Scenario: To show how resource exhaustion can be tackled by enforcing a global performance optimization goal. The limits of the system are explored using an experiment with a large number of virtual nodes and many management information flows. Scalability refers to the ability of the IS to handle growing networks elements and usage in a graceful manner and its ability to be enlarged to accommodate that growth. Stability refers to the degree to which IS must work/operate in a changing environment.

We stress test our S/W infrastructure with large topologies (up to 500 virtual nodes). The main goal here is to investigate its behaviour in terms of scalability and stability. As is shown in Figure 5-19, Figure 5-20 and Figure 5-21, large scales can be reached. Figure 5-19 highlights how IS CPU load increases with the topology size. Since the number of management information flows remains the same, there is no impact on the memory consumption (Figure 5-20). The next figures (Figure 5-22, Figure 5-23 and Figure 5-24) show how IS can trade an increased jitter in response-time for a slight increase in the average response time in the case of a large scale topology and gradual resource exhaustion. In this example, we enforced a global performance goal change that switches the communication method from Pull to Direct Communication. This strategy can be associated with a control loop that detects and tackles systematic stability problems.

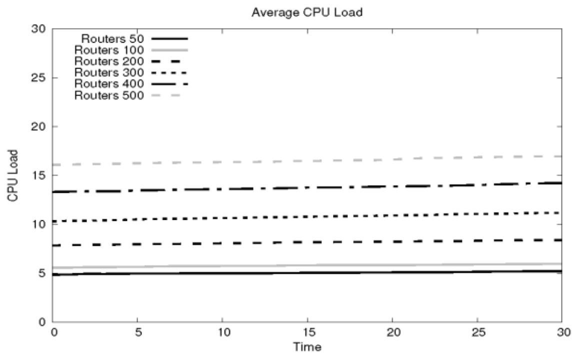

Figure 5-19: IS CPU Load (Direct Communications)

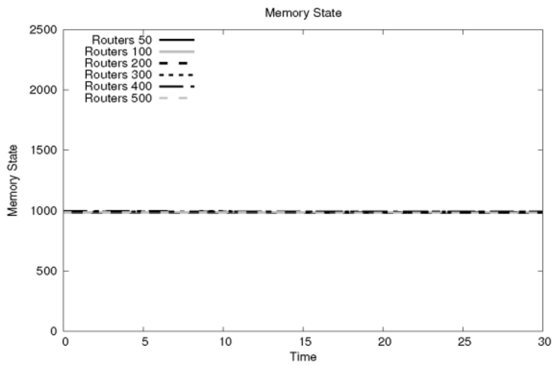

Figure 5-20: IS Memory Consumption (Direct Communications)

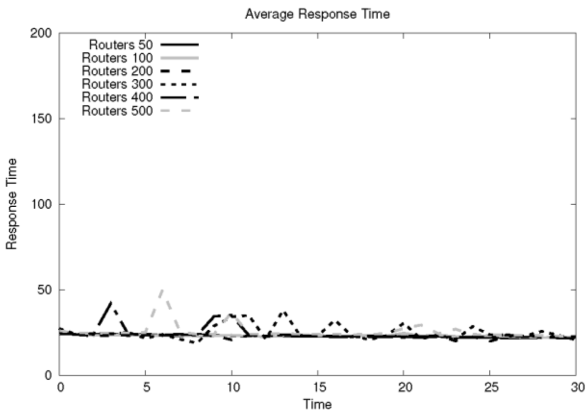

Figure 5-21: IS Average Response Time (Direct Communications)

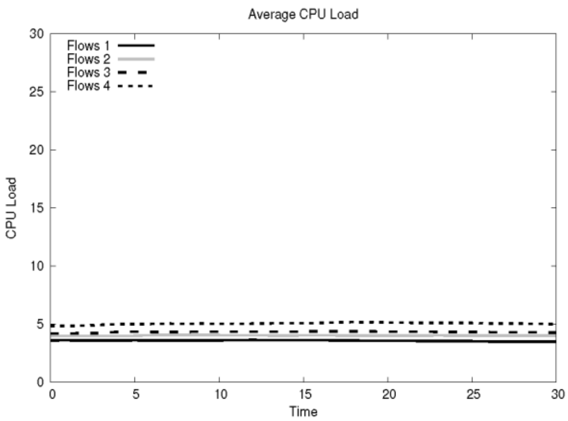

Figure 5-22: IS CPU Load (Handling Jitter)

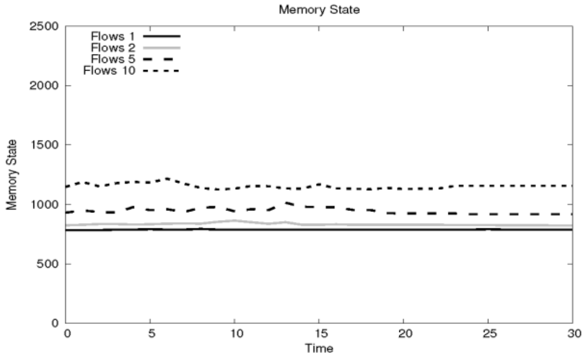

Figure 5-23: IS Memory Consumption (Handling Jitter)

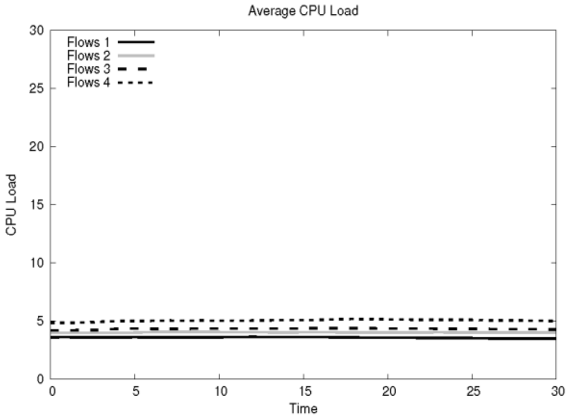

Figure 5-24: IS Average Response Time (Handling Jitter)

6 Concluding Remarks and Lessons Learned

6.1 Concluding Remarks

This document presented the results of integration, validation and performance evaluations of the mechanisms implemented within the DOLFIN project.

It took into account all of the implementation results of the components developed in WP3 and WP4 and the testbed description provided in D5.1. This deliverable elaborated the validation and evaluation of the DOLFIN prototype representing the final outcome of task 5.3. It provided a clear description of the performances that the DOLFIN prototype can offer as a solution for the efficient energy management of Data Centres.

The experiments carried out on the integrated DOLFIN platform followed the three scenarios defined in WP2. We took into account the DOLFIN DC categories as particular context for the experiments and we also have addressed the challenges identified by the project.

This document presented how the DOLFIN platform behaves in the proposed scenarios using the optimization policies implemented in the eCOP and SDC components. The analysis carried out highlighted which is the estimation of the amount of energy saved and consequently which the estimation on the revenue's benefits in relation to the utilization of DC resources, according to the evaluation criteria identified in the project.

The evaluation activities and results have contributed to the consortium-defined exploitation strategy documented in deliverable D6.4 as well as for each partner's specific exploitation plans.

6.2 Lessons Learned

The DOLFIN project completed its core development and evaluation activities, briefly overviewed in this deliverable. The project allowed for the successful and harmonic cooperation of nine partners from six different countries (Italy, Spain, Greece, United Kingdom, Romania and Poland) of different ICT developments and concepts around DC S/W platforms, operations and cloud computing management. However, the coordinated effort of the consortium enabled the development of a novel energy–aware DC cloud platform leading to significant findings and lessons learned, some of which were already presented as evaluation results in this deliverable. The most significant ones are discussed below:

1. **IT load (VMs) consolidation can assist in greatly reducing the energy consumption of a DC**. In general, DCs tend to be highly underutilized; usually 30% of DC servers are in very low utilization or even comatose state and are mainly used to accommodate peak loads that the currently active IT equipment cannot handle successfully and under specific Quality of Service constraints [32]. Shutting down these servers and move their IT loads to other already operating servers can substantially decrease the overall DC consumption. However, from a pure business perspective and in operational environments, one has to consider that peak loads are not rare and they actually matter when it comes to cloud computing services reliability and firm reputation. In order to combine the operational needs for increased services resilience while achieving energy savings out of IT load consolidation, the DC Operators should decide on a mix balancing this trade-off. Our DOLFIN experience indicates that most servers are able to enter a sleep state in a few seconds after a relevant command is sent. The wakeup procedure is also very fast, also being measured in seconds. Note that server shutdown is not recommended as sleeping and wakeup times are much larger; setting a server at sleeping state causes it to achieve minimal power dissipation (e.g. in the case of the PSNC testbed the power dissipation at sleeping state of the servers detailed in 4.3.3.3 is just 12W/server) while keeping wakeup times at acceptable levels.

   Recommendation**:** In such a framework, keeping a percentage of the servers (e.g. 10%) at standby/inactive state to handle the very first waves of the peak loads, allowing the rest of the servers to successively wakeup from the sleeping state and handle the rest of the load, could prove to be a valid solution to the performance vs. energy efficiency dilemma, the percentage of standby servers being the key towards managing this trade-off. Advanced energy and correlated user-behavioural predictive analytics are keys towards identifying in an automated manner to adjust the percentage of active vs. sleeping servers, increasing the relevant ratio when an increase in the IT load is expected and lowering it in the opposite case.

2. **The energy benefits change with the DC size**. As apparent in the general DOLFIN evaluation section, the effectiveness of the application of DOLFIN depends on many parameters, DC size being one of them. Although DOLFIN can deliver significant benefits in smaller DCs, its effect can be enabled with a larger impact when applied to larger DC scales. This is due to the fact that larger DCs have a greater chance of maintaining unused equipment for peak-load service provisioning and also that larger DCs offer more flexibility as to the possible configurations that may be exploited in the context of optimal resources management exploration.

3. **Keeping in track with standards does matter**. In the course of conducting our standardization and dissemination activities, particularly in industry-oriented contexts such as EMAS, we confirmed that keeping in track with standards greatly improve the chances for real (post-piloting) adoption both from the consortium partners and the outside world. The DOLFIN-developed DCO Brokers are able to interface all core OpenStack monitoring and control modules, whereas adaptation to VMWare is considered equally (or even more) important. As VMWare is the de-facto cloud management platform of most cloud providers (OpenStack is the emerging one but still not as mature as its great opponent), supporting it could result in a better exploitation potential, overall. In the same manner, the DCO Brokers are able to interface IPMI and SNMP interfaces for accessing and controlling IT and non-IT equipment, something that proved to be valuable in the integration phase, as integrating control of the equipment was straightforward and trustworthy in terms of results.

4. **IT load migration execution time is unpredictable**. Although VM migration is natively supported by all available virtualization managers and cloud management platforms, its performance dramatically depends on the nature of the load supported by the various VMs. VMs hosting CPU-intensive tasks can be migrated to other physical nodes in the same DC in less than five seconds when shared VM storage is assumed, the respective time in the case of non-shared storage can be many times larger (depending on the ratio of hard disk versus total physical memory). In the same context, applications with excess RAM or IO usage tend to be significantly less time-efficient in migration terms, as the continuous updates in the contents of the memory/disks successively invoke incremental memory/disks snapshots that, granted high rates of writing, can result in unacceptable execution timings in operational terms; while performing the DOLFIN performance evaluations, we faced a situation of a VM that took over two hours to get migrated due to excess disk activity.

   Recommendation: Grouping the IT loads based on their core application activity can help towards identifying which VMs are good candidates for migration and which are not. This will decisively affect the time of the execution of the optimization plans.

5. **Cross-DC IT load relocation is time-consuming, harder to achieve but also very rewarding**. In the context of the Smart Grid integration scenario (4.3.3) we demonstrated that cross-DC IT load relocation is possible. As documented in the relative section, the testbed setup was using a 50km long optical fibre cable to interconnect the two racks representing the different DCs. However, even in this nearly optimal case (short distance for two DCs and optimal, laboratory conditions for the fibres from a communications perspective), each VM relocation took on average 1.5 minutes to complete (the average VM characteristics were 1.5GB of RAM and approximately 10GB of disk storage). These times can be considered acceptable in delay-tolerant workloads and SLAs but also completely unacceptable in other cases. Moreover, VM traceability is an important issue as when the VMs change DC their UUIDs change; DOLFIN solved this problem by using a distributed DB dedicating for storing cross-DC information, namely the DOLFIN Info DB. This is one of the reasons that make cross-DC IT load relocation harder to achieve in practice, as most DC Operators are not willing to share such information. On the other side, as apparent from paragraph 4.3.3.7, cross-DC IT load relocation can result in very significant energy savings and improvement of the overall energy mix, also enabling wider green SLA adoption by the various DCs that are already powered by green sources and would like to have IT load hosting alternatives. Last, IT load migration was proved to be arguably the best way of achieving compliance to extreme DSO DR requests in case of Smart Grid emergency situations.

   Recommendation: In the context of Smart Cities and particularly Smart Grids, IT load relocation can bear very significant merits. However, the lack of a standardized solution as to SLA monitoring, VMs traceability as well as confidentiality factors hinder the relevant adoption by the DCs, hence depriving Smart Grids from a great grid-balancing option. The relevant policy making bodies should consider building a framework that would allow such options while preserving confidentiality, security, privacy and guaranteeing data non-repudiation.

6. **Liaisons with other projects are more valuable than though**. Although the project partners share very different experiences and very diverse areas of expertise, it happens that targeted discussions tend to annihilate the different point of views of the consortium as a whole, limiting the possibility to explore better options to solve a particular problem. Although staying focused is definitely an arrow in a project's quiver towards success, this can

also result in a loss of opportunities. We found our liaison with other projects of the same (e.g. GEYSER) and different scope (e.g. FINESCE or 5G projects) to be enlightening in many ways, from optimization procedures to integration processes, also spanning from understanding how Smart Grid balancing can be achieved to how virtual network functions can be abstracted in a core computational environment. Also, our experience in the context of the DC Cluster indicates that joint EC projects effort can have a much more practical impact on the various standardization bodies.

Recommendation: Clustering projects are very beneficial to all involved projects. We firmly believe that they should be further encouraged by the EC as this would result in much higher impacts achieved.

7. **Modular architectures are more resilient and sustainable than monolithic ones**. The implementation of DOLFIN was based on a decentralized, modular architecture that included components of completely different scope, using different technologies and sharing very diverse functional and non-functional requirements. As the project evolved, several components were subject to severe re-design processes either to support newly identified functionality or to circumvent practical implementation and operational limitations. The adopted modular architecture allowed us to perform such changes without affecting the interoperability with other components; keeping the interfaces intact would guarantee that everything would work well after the refactoring. In the same course, a DC Operator could easily replace/extend a DOLFIN component reference implementation to gain added value or differentiated services; an example could be an extension of the Optimizer to apply different kinds of optimization algorithms.

8. **Dynamic self-documentation frameworks and automated VM generation speed up development and integration efforts**. The vast majority of the components delivered in the context of DOLFIN expose dynamic self-documentation interfaces that allow for quick and easy experimentation and testing. It was agreed by all partners that maintaining such infrastructures significantly speed-up development efforts, minimize the required communication activities and result in achieving faster and fault-preventing integration. In such a framework, the existence of such documentation services in the DOLFIN components could be appreciated by third parties that would like to re-use part of the DOLFIN components (either distinctively or in groups) for own purposes.

   In a similar context, the open-source, auto-generated, auto-configurable and auto-deploying DOLFIN-powered VM greatly assisted us in achieving seamless integration and testing as all developing parties could have access to a valid, up-to-date, local DOLFIN installation to perform different types of tests and check different scenarios without the fear that an accidental fault or erroneous change could result in an integration disaster.

9. **Translating every activity in a single context is critical for the design of energy efficiency platforms**. When designing energy-efficiency solutions, it is critical that all computational processes, informed decisions and applied actions are based on the translation of the data and the control impact in terms of energy. During the project development and evolution, the focus of the consortium shifted from a pure ICT to a mixed energy-efficiency-through-ICT perspective; the extensive use of energy models to accurately calculate the energy consumption per VM decisively contributed to achieving DOLFIN's consolidated view of IT loads as a form of energy consumption rather than applications running on VMs. Based on that homogeneous view, the various DOLFIN components were able to "speak the same

language" and operate towards achieving fine-grained energy consumption control, rather than VM mobility.

7 References

[1] DOLFIN, "Deliverable D5.1: Testbed description and testing scenarios," 2016.

[2] DOLFIN, "Deliverable D6.4: Final Exploitation Plan," 2016.

[3] Louis Columbus, "Roundup Of Cloud Computing Forecasts And Market Estimates, 2016," Forbes, [Online]. Available: http://www.forbes.com/sites/louiscolumbus/2016/03/13/roundup-of-cloud-computing-forecasts-and-market-estimates-2016/.

[4] NRDC, "Data Center Efficiency Assessment," August 2014. [Online]. Available: https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf.

[5] ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY, "United States Data Center Energy Usage Report," June 2016. [Online]. Available: http://eta.lbl.gov/sites/all/files/lbnl-1005775_v2.pdf.

[6] NRDC, "The carbon emissions of server computing for small- to medium-sized organizations," Ocbober 2012. [Online]. Available: https://www.nrdc.org/sites/default/files/NRDC_WSP_Cloud_Computing_White_Paper.pdf.

[7] DOLFIN, "Deliverable D2.2: DOLFIN requirements and system architecture," 2014.

[8] Texas Advanced Computing Center, "Hikari: Sustainable Supercomputing," Texas Advanced Computing Center, [Online]. Available: https://www.tacc.utexas.edu/systems/hikari.

[9] Texas Advanced Computing Center, "New hikari supercomputer starts solar hvdc," Texas Advanced Computing Center, [Online]. Available: https://www.tacc.utexas.edu/-/new-hikari-supercomputer-starts-solar-hvdc.

[10] DOLFIN, "Deliverable D4.1: Synergetic Data Centres for energy efficiency (Design)," 2015.

[11] APC, "Calculating Total Cooling Requirements for Data Centers," [Online]. Available: http://apcmedia.com/salestools/nran-5te6he/nran-5te6he_r3_en.pdf. [Accessed 07 2016].

[12] DOLFIN, "Deliverable D3.1: Data Centre energy consumption optimization platform (eCOP) Design," 2015.

[13] HashiCorp, "Vagrant by HashiCorp," [Online]. Available: https://www.vagrantup.com/.

[14] IBM, "Agile DevOps: Infrastructure automation," [Online]. Available: http://www.ibm.com/developerworks/library/a-devops2/.

[15] HashiCorp, "Vagrant - Providers," [Online]. Available: https://www.vagrantup.com/docs/providers/.

[16] Github, "Vagrant Openstack Provider," [Online]. Available: https://github.com/ggiamarchi/vagrant-openstack-provider/.

[17] DOLFIN, "DOLFIN source code repository - Vagrant VM," [Online]. Available: https://stash.i2cat.net/projects/DOL/repos/y2_demo1_vm/browse.

[18] Supervisor, "Supervisor: A Process Control System," [Online]. Available: http://supervisord.org/.

[19] © Intel Corporation, "Intelligent Platform Management Interface (IPMI)," [Online]. Available: http://www.intel.com/content/www/us/en/servers/ipmi/ipmi-home.html.

[20] © Intel Corporation, "IPMI Adopters List," [Online]. Available: http://www.intel.com/content/www/us/en/servers/ipmi/ipmi-adopters-list.html.

[21] Openstack, "Kwapi," [Online]. Available: https://launchpad.net/kwapi.

[22] OpenStack, "Telemetry measurements," [Online]. Available: http://docs.openstack.org/admin-guide/telemetry-measurements.html.

[23] C. Kaner, "An Introduction to Scenario Testing," June 2003. [Online]. Available: http://www.kaner.com/pdfs/ScenarioIntroVer4.pdf. [Accessed September 2015].

[24] "Openstack," [Online]. Available: https://www.openstack.org/.

[25] DOLFIN, “Deliverable D3.3: ICT Performance and Energy Supervisor component (Implementation),” 2016.

[26] DOLFIN, “Deliverable D3.4: Energy Efficiency Policy Maker and Actuator component (Implementation),” 2016.

[27] “openADR Alliance,” [Online]. Available: http://www.openadr.org/.

[28] “Mirantis,” [Online]. Available: https://www.mirantis.com/.

[29] “OpenStack Fuel,” [Online]. Available: https://wiki.openstack.org/wiki/Fuel.

[30] “OpenStack Configuration Reference,” [Online]. Available: http://docs.openstack.org/kilo/config-reference/content/.

[31] DOLFIN, “Evaluation Framework source code,” [Online]. Available: https://stash.i2cat.net/scm/dol/evaluation_framework.git.

[32] B. Kepes, “30% Of Servers Are Sitting "Comatose" According To Research,” Forbes, 3 June 2015. [Online]. Available: http://www.forbes.com/sites/benkepes/2015/06/03/30-of-servers-are-sitting-comatose-according-to-research. [Accessed 13 October 2016].

[33] DOLFIN, “Deliverable D2.1: Business scenarios and use case analysis,” 2014.