dolfin

The Green Data Center: *solutions for* Energy-Efficient DCs

www.dolfin-fp7.eu

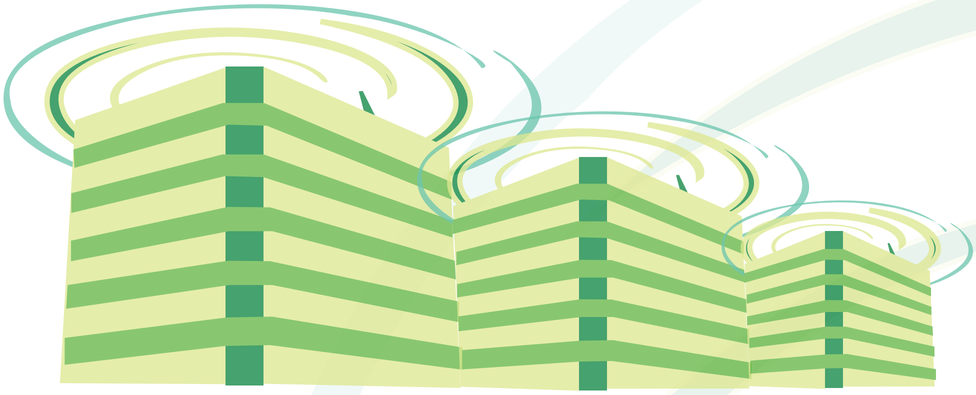

1. Introduction

During the last few years, the energy sector has faced two very important challenges:

a) **increase of active Data Centres (DC),** with subsequent increase of energy demand/consumption and sub-optimal energy management. This directly affects the energy footprint of DCs and environmental conditions. It is well known that the average server utilization in DCs is low, often below 30% of the maximum server load and only 10% in case of facilities that provide interactive services. The over-dimensioning of DCs and the increased number of underutilised servers has significantly increased energy demand and consumption, and a huge amount of energy is also consumed for the cooling of DC servers. To lower this waste of energy, DC containment strategies (both hot and cold aisle) are widely regarded as the starting point for energy efficiency.

b) **instability of the Smart Grids and their difficulty to follow the electricity demand-response model.** As Europe shifts away from fossil fuels towards integrating larger shares of renewable energy sources, electricity is becoming an even more important energy vector. Smart Grids are very high on the agenda of the European energy and ICT sector; however, difficulty in following the electricity demand-response model still remains severe.

DOLFIN tackles these problems and significantly contributes towards improving the energy efficiency of DCs and the stabilisation of Smart Grids, following a holistic approach, across networks of DCs and Smart Grids. DOLFIN models, monitors and measures energy consumption and enables seamless, autonomic migration of Virtual Machines (VMs) between servers of the same DC or distribution across a group of energy-conscious, synergetic DCs, aiming to i) optimize the overall energy consumption 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 [1].

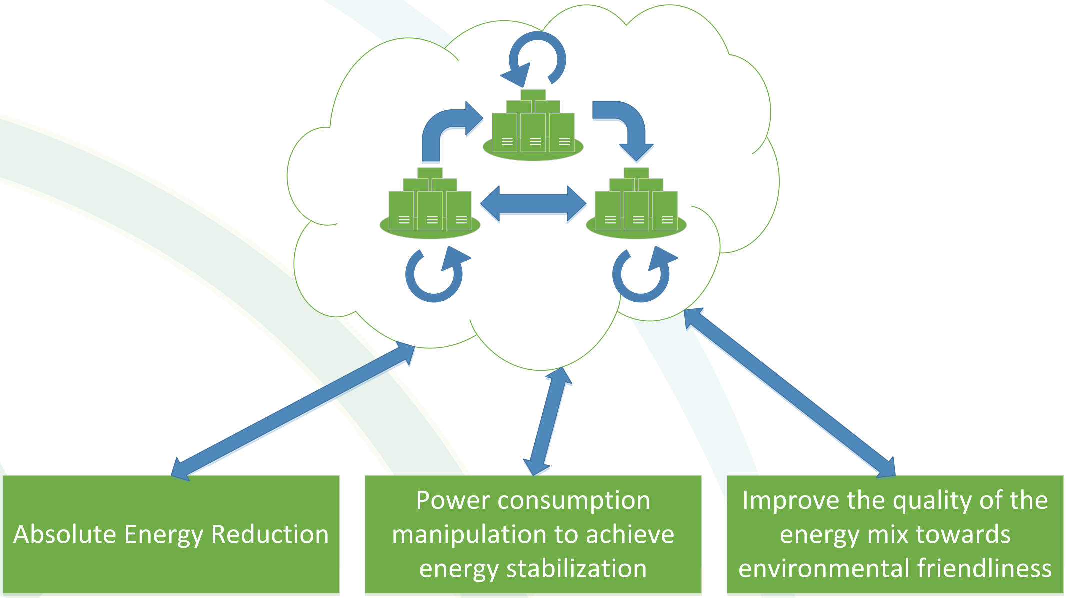

Figure 1: The main classes of the DOLFIN business scenarios.

Two main optimization and control loops governing DC operation: one related to self-optimization (e.g. each DC optimizes its own state without considering other synergetic DCs in a common DOLFIN framework) and one related to inter-DC operation (e.g. each DC is operated in a coordinated manner, the optimization process referring to the DCs as a whole).

Considering this generic view, the DOLFIN developments have been based on a modular approach, separating the local (intra-) DC optimization control processes from the inter-DC ones, only allowing a link between them to be exposed through a set of well-defined interfaces. Morever, even in the case of intra- and inter-DC optimization, all functionalities have been implemented in a modular way, so that re-use of components and/or replacement with other implementations (other than the DOLFIN-reference ones) is possible,increasing the flexibility of the DOLFIN solution.

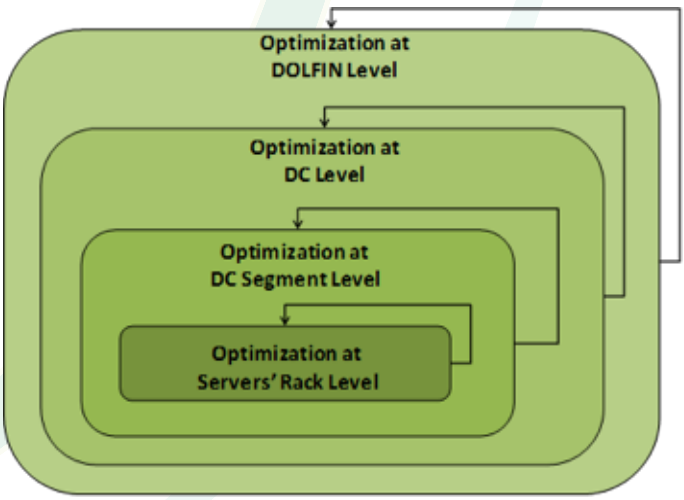

Figure 2: DOLFIN multi-layered optimization approach.

2.DOLFIN Architecture

DOLFIN implements followed a three levels hierarchical approach for optimizing energy consumption:

Level 1 - Optimization inside a Data Centre: Optimizing the energy consumption within the limits of a single DC, based on system virtualization and the optimal distribution of VMs. This was coupled with the dynamic adaptation of active and stand-by servers and the load optimization per active server. DOLFIN employed a new monitoring framework to measure the energy consumption per server module/ networking component and activate low-power states on devices. DOLFIN complemented these energy management mechanisms by prototyping a system for heat absorption from the CPUs of the physical servers, through the use of cold water heat management systems.

Level 2 - Group optimization of Synergetic Data Centres: Optimizing the cumulative energy consumption in a group of DCs, based on optimal distribution of VMs across all of the servers that belong in the group of DCs using policy-based methods. DOLFIN measured the energy consumption on the DC level and deployed policies and solutions capable to achieve a decrease in the aggregate power consumption across the whole group of synergetic DCs.

Level 3 - Optimization at Smart City/Smart Grid Stabilization: Optimizing the energy consumption at the smart city level and stabilizing the Smart Grid, based on distribution of VMs across a group of DCs, following an electricity demand-response approach. In order to stabilize the Smart City Energy consumption, DOLFIN dynamically changed the percentage of active versus standby servers per DC and optimized the load per active server.

To enable this multi-level optimization continuous monitoring solutions should be employed at the finest possible level of granularity, interfacing with already existing S/W and H/W infrastructures. In the same framework, control over heterogeneous S/W and H/W infrastructures should be possible. To this end, all DOLFIN components interact with the actual DC infrastructure through a set of well-defined brokering mechanisms that pertain to monitoring and control of the physical and virtual infrastructure (including Hypervisors, Monitoring mechanisms and IT- or non-IT appliances automation systems).

The DOLFIN architecture [2] is consolidated around two specialized sub-systems:

Energy Consumption Optimisation Platform (eCOP), which represents the core platform in DOLFIN as previously discussed, and is where the relevant actions are performed to manage and optimize the energy consumption at Data Centre level.

Energy-conscious Synergetic DCs (SDC), which provides a dynamic, service-effective and energy-efficient allocation of demands, across a distributed network of co-operating DCs. In addition, the SDC provides the control modules for the integration with the Smart Grid environment and that are responsible for the energy stabilisation through the interconnection with the smart grid network, responding to the demand side management directives of the latter.

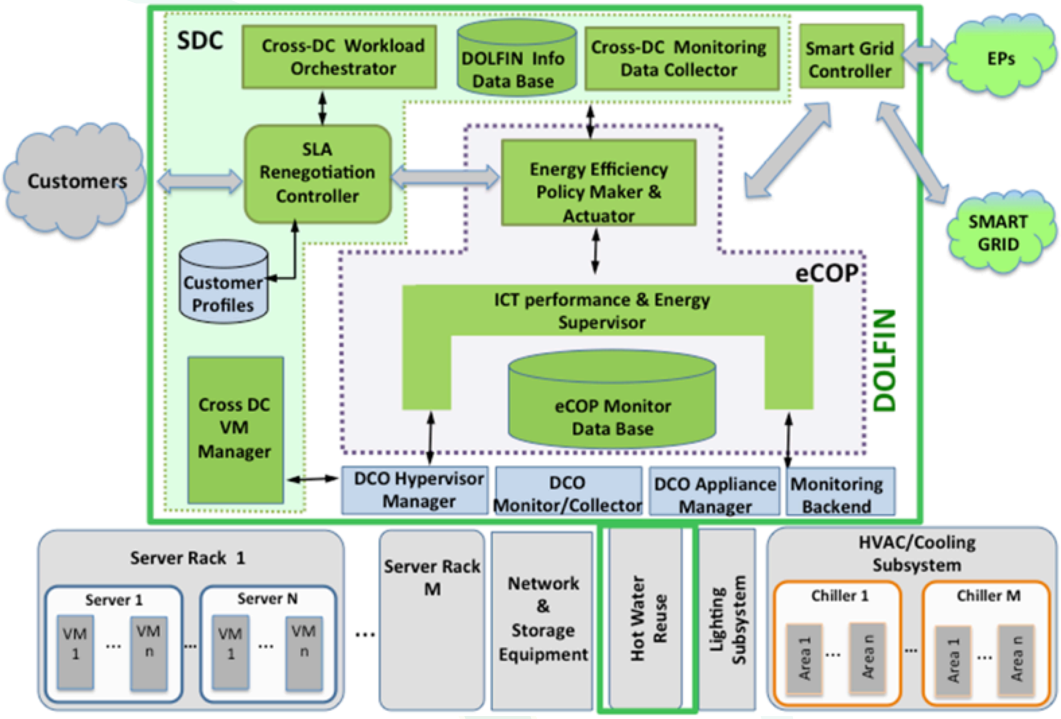

Figure 3: The functional architecture of DOLFIN [2].

2.1.Data Centre Energy Consumption Optimization Platform (eCOP)

The eCOP platform [3] is the core set of functional modules in the DOLFIN architecture that grants a DOLFIN-enabled DC with advanced monitoring, self-evaluation and self-optimisation capabilities. The four pillars of eCOP functionality are summarized as follows:

- Interfaces legacy/heterogeneous DC sub-systems (including cloud management platforms and building management systems) by means of specific adapters to collect all necessary information for the optimization process.

- Offers an engine able to process raw data and calculate appropriate metrics and KPIs that will be used throughout the whole the optimization process. This engine resides in the ICT Performance and Energy Supervisor module.
  Implements mechanisms to store raw data, measures and calculated metrics into the eCOP DB, and provides also efficient and abstracted interfaces to support DB operations.

- Sets the Energy Efficiency Policy Maker & Actuator module into action to perform the optimization lying in the core of the DOLFIN project scope. The module activation is based on pre-defined criteria and aggregated information from other modules (e.g. ICT Performance and Energy Supervisor, Smart Grid Controller, Cross-DC orchestrator, etc.).

The eCOP interfaces to legacy control and orchestration modules of the DC (DCO modules) grouped in three areas of functions:

- **Hypervisor,** which is responsible for managing new ICT hardware and VMs configuration;
- **Appliance**, which acts on the non-ICT DC infrastructure (e.g. HVAC);
- **Infrastructure**, which interfaces with both the ICT and the non-ICT DC infrastructure to collect all operational and energy related information to be stored in the eCOP.

The current implementation of the **DCO Hypervisor Broker and the DCO Monitor/Collector** have, both, focused on two platforms, that is **OpenStack** [4] and **VLSP** [5], the software-defined DC infrastructure provided by UCL.

Three main components implement the eCOP functional perspective:

www.dolfin-fp7.eu

1. The **ICT Performance and Energy Supervisor module** has the responsibility to interact with underlying legacy DC subsystems and collect relevant information to evaluate metrics and KPIs. The module implements mechanisms to efficiently retrieve the data from various sub-systems and distribute such information to appropriate consumers. Moreover, the ICT Performance and Energy Supervisor undertakes data analysis and representation, useful to other eCOP module for their optimization activities.

2. The **Energy Efficiency Policy Maker and Actuator** implements the intelligence needed for the derivation and application of DC optimization policies. It is responsible for the application of the energy optimization procedures based on particular predefined criteria and conditioned by the inputs (requests) provided by other DOLFIN components, for example requests from the Smart Grids environment, from federated DCs, etc.

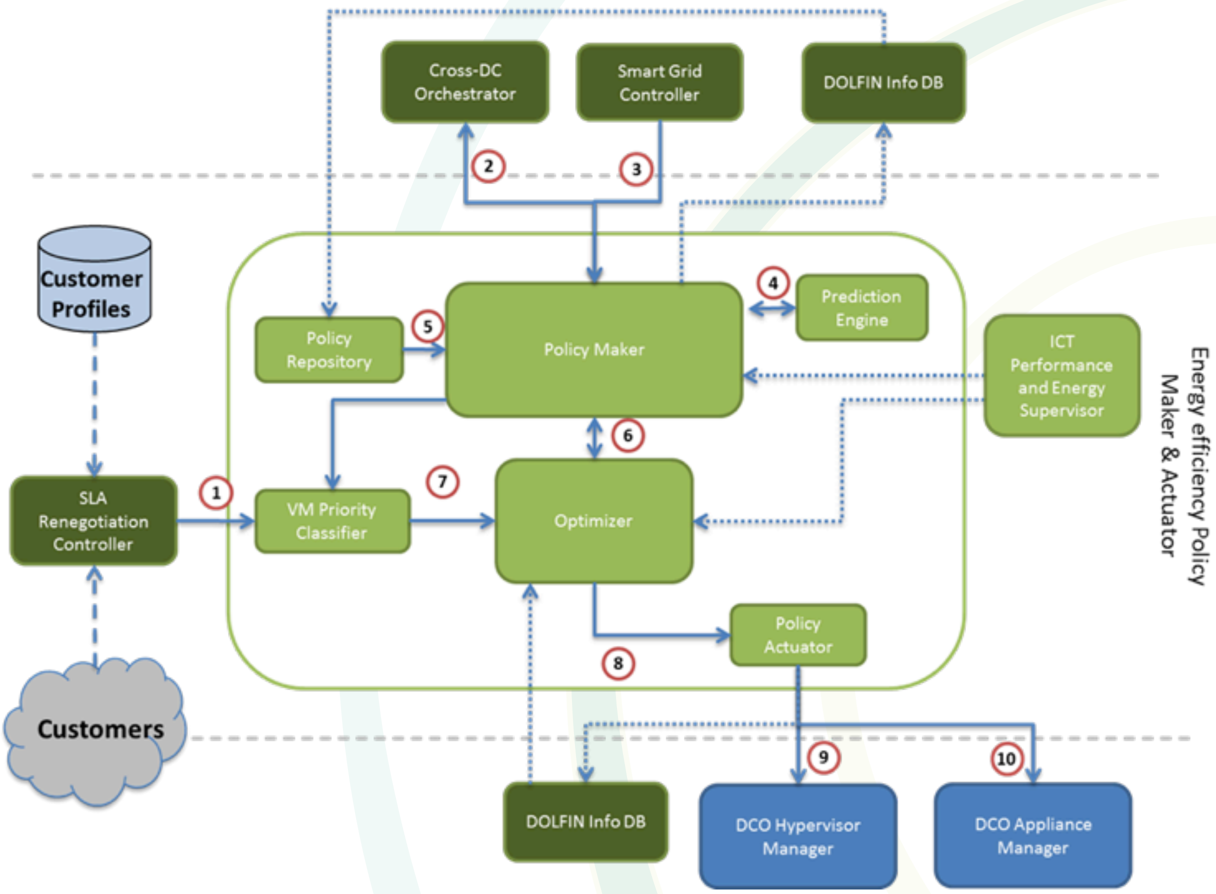

Fig.4: The Energy Policy Maker & Actuator.

3. The **eCOP Monitor Database** groups all the elements involved in the storage functionalities for measures, KPIs, DC assessment and status, etc. The eCOP Database not only implements pure database functions, but introduces specialized modules to efficiently interact with that data store and produce a time series aggregation..

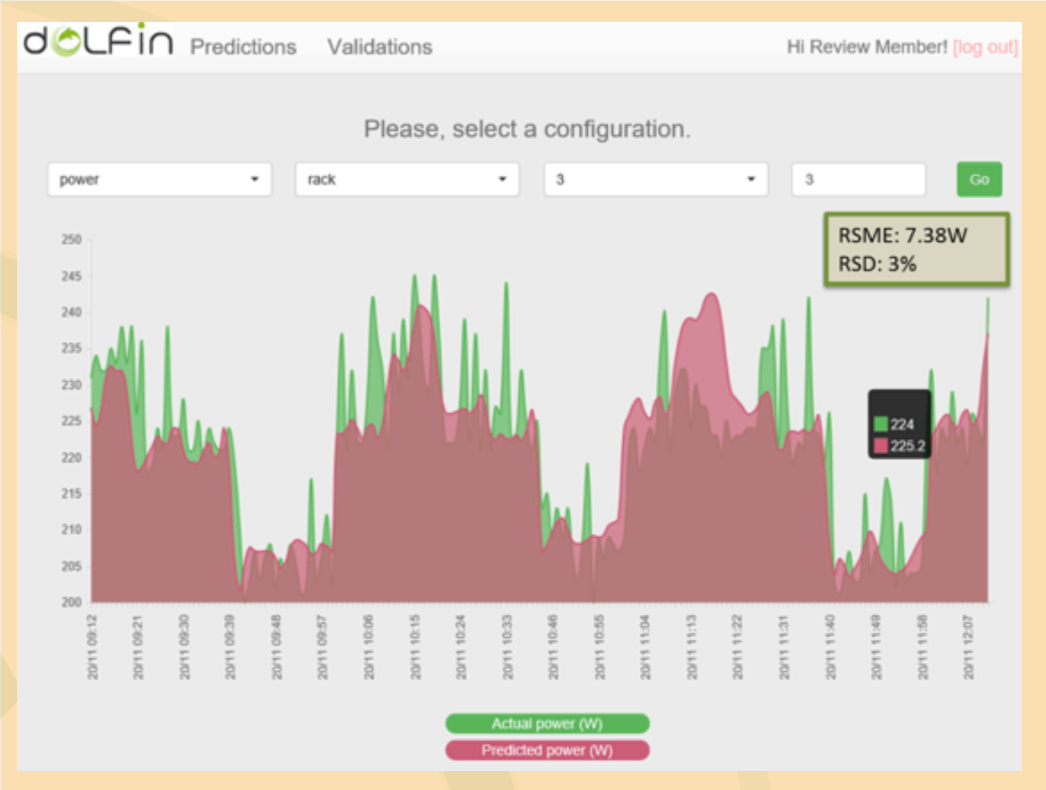

Fig. 5: Actual power demand vs forecast by the Prediction Engine in DOLFIN eCOP

2.2.Water cooling server module

The water cooling server module developed within DOLFIN, follows a three-level hierarchical approach:

- at the lower-level (3rd level), a water-loop is utilized per server for conveying the heat of the high power components (CPUs, GPUs, RAM, hard-disk) out of the server
- at the second level, a water loop is utilized per rack for interconnecting all the
- 3rd level loops within the rack and thus enabling the heat energy of all the local loops to escape out of the rack.
  last, a top level loop is utilized for interconnecting the 2nd level water-loops of all the racks and thus enabling to convey the heat energy of all the racks out of the computer room.

Furthermore, the top-level loop is interconnected to the conventional HVAC system of the DC and thus the heat energy is re-used for conventional heating purposes (i.e. space heating). The following figure depicts the aforementioned approach.

In cases of very large and very crowded computer rooms, this hierarchical approach could be expanded to include an extra level of water-loops for interconnecting groups of racks together and thus enabling better and more fair heat dispersal.

As can be easily deduced from the above, the role of the three loops is related to the different steps of heat absorption from the servers and the heat transfer from them to the HVAC system;

Figure 6: The water cooling server module design

- 1st level liquid loop: The role of the top level loop is to interconnect the liquid loops of all the racks of a computer room into a single loop and thus enabling to handle all the generated heat as a whole. The heat exchangers utilized for the interconnection of the distinct liquid-loops are plate-type heat-exchangers.
- 2nd level liquid loops: The second level liquid-loops are at the rack level, this means that their role is to interconnect all 3rd level liquid loops that are housed inside the rack into a single loop that in its turn is interconnected into the top level loop. The interconnection of the 3rd level liquid loops to the 2nd level loop is realized by simple T-joins.
- 3rd level liquid loops: Finally, the 3rd level liquid loops are at the server level. Their use is to convey the heat generated by high power server components out of the server and into the 2nd level liquid loop of the respective rack.

2.3.Synergetic Data Centres for energy efficiency subsystem (SDC)

The SDC subsystem is responsible for granting DOLFIN-enabled DCs to create federations of DCs able to exchange performance and load-related information so that energy efficiency awareness can be achieved at a global, federated level. Further, being logically situated above the eCOP subsystem, the SDC is responsible for allowing the interaction of the DC with its environment, including Smart Cities, the Smart Grid and Energy Providers (EPs) and the DC Customers

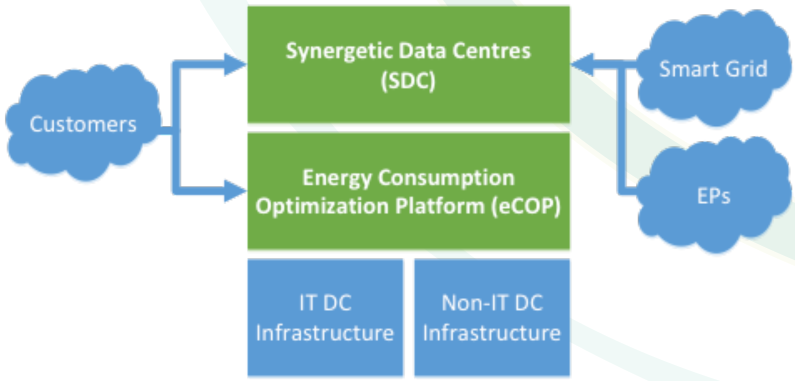

Figure 7: The logical placement of the SDC in the general DOLFIN context (green).

In this context and with reference to Figure 0 1, the SDC has been implemented as a single subsystem composed of six different components, namely:

- The Cross-DC Workload Orchestrator;
- The Cross-DC VM Manager;
- The Cross-DC Monitoring Data Collector;
- The Smart Grid Controller;
- The SLA Renegotiation Controller (also
- aware of the customer profiles);
- The DOLFIN Info Database.

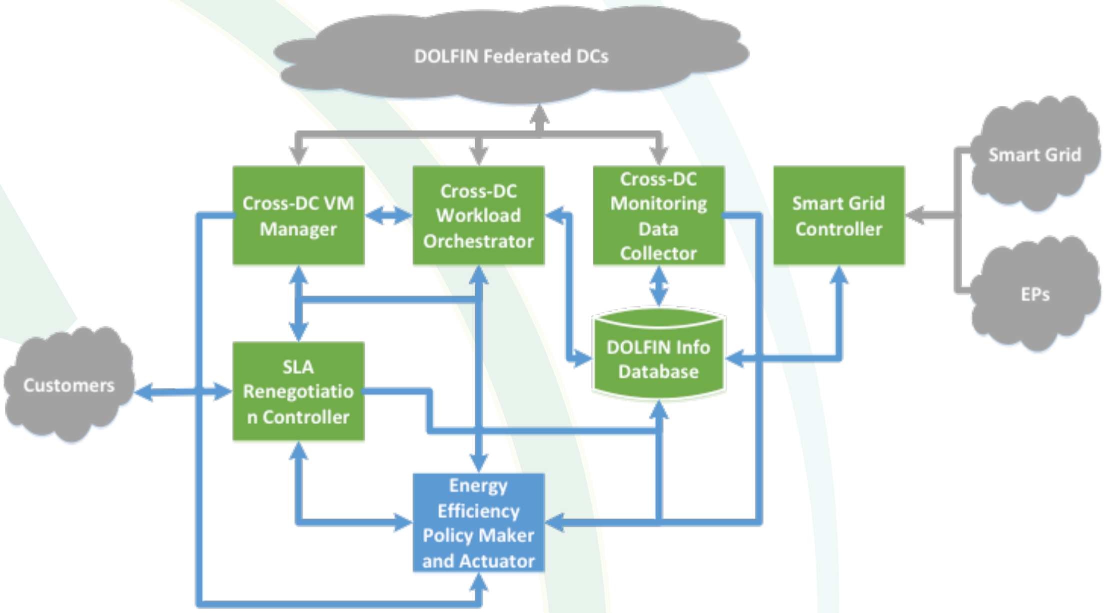

Figure 8: SDC functional architecture (green).

The generic SDC architecture is presented in Figure 0 4, with emphasis on the interaction of the SDC components with the outside world (coloured in grey) and the Energy Efficiency Policy Maker and Actuator (in blue colour)

The component implementing the core functionality and control flow handling in the SDC plane is the **Cross-DC Workload Orchestrator**. It is a distributed software element that gets the SDC decisions and does most of the different types of resource optimisation including energy optimisation and management of their trade-offs in a federated DC optimization scenario. It handles DC migration message requests by performing negotiation with other CDC Workload Orchestrators located in distributed, DOLFIN-federated Data Centres in order to find one willing to receive the VM(s) that need to be relocated.

3.Performances

Three basic scenarios were identified for validating the entirety of the integrated DOLFIN functionality [6]:

Scenario 1 optimize operation at intra-DC level, lowering the overall DC energy consumption;
Scenario 2 SLA-awareness to trigger a new optimization process to force the SLA respect.
Scenario 3 notification from the Smart Grid Controller about a change in the energy prices in the next hours can trigger a new optimization plan that results in actually relocating some VMs (hence IT and non-IT load) to another DOLFIN federated DC.

All these scenarios combined, present the intra-DC optimization, flexible SLA handling and inter-DC optimization under Smart Grid limitations capabilities of DOLFIN.
The three tests were executed in the three different infrastructures, the first scenario being operated in WIND testbed, the second in Interoute and the third one in PSNC.

All tests were successful, indicating that DOLFIN overall functionality is concrete and satisfies the requirements for boosting energy efficiency, flexible SLAs monitoring and handling as well as successfully integrating Smart Grid Operators notifications and federated (inter-) DCs orchestration.
We were able to test more than 600 different DC topologies and more than 12,000 different DC configurations in total.

For small-size urban DCs, holding 2 – 5 racks with 5 – 25 servers, respectively energy savings can reach 75% when the Policy is set to optimize against **energy efficiency in case of low utilization.** With increased DC utilization there are reduced gains, though still significant savings of approximately 20%.

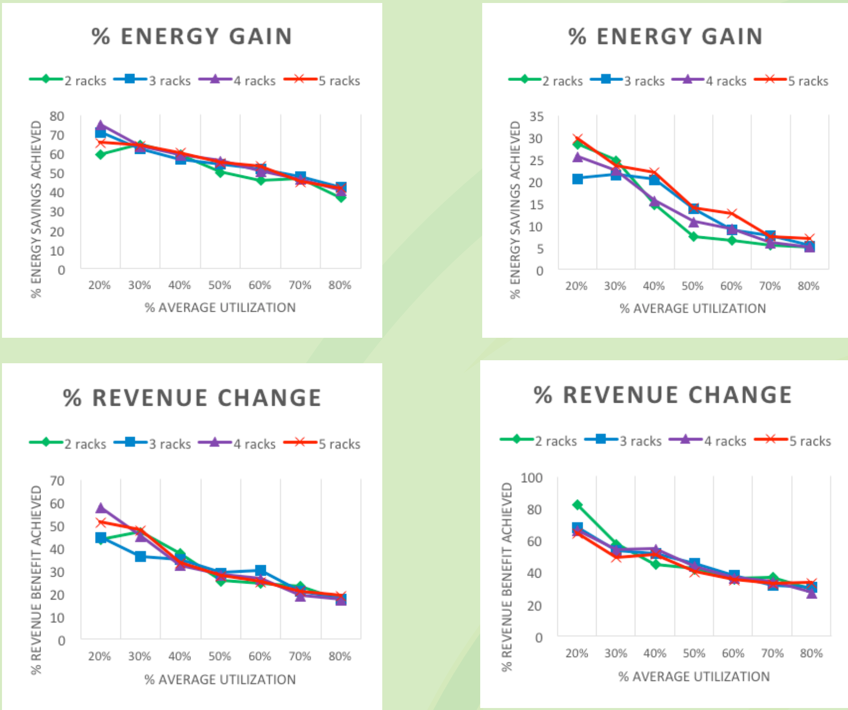

Figure 9 : DOLFIN performance on small urban DCs when the DC policy is set to optimize against energy efficiency (left column) or performance (right column).

In the case of medium-sized DCs operating against 20 – 50 racks, hosting 100 – 250 physical servers and 10% of the DC infrastructure powered by green sources, **the performance of DOLFIN is able to provide energy-efficiency and revenue increase of over 80% in certain cases (very low utilization).** Even in the case of highly-utilized DCs, the energy and revenue benefits acquired through DOLFIN are very significant, reaching 60% in the case energy or performance policy is applied, respectively. In general, the effect of DOLFIN in partially green-powered medium-sized DCs is shown to be greater when compared to small ones.

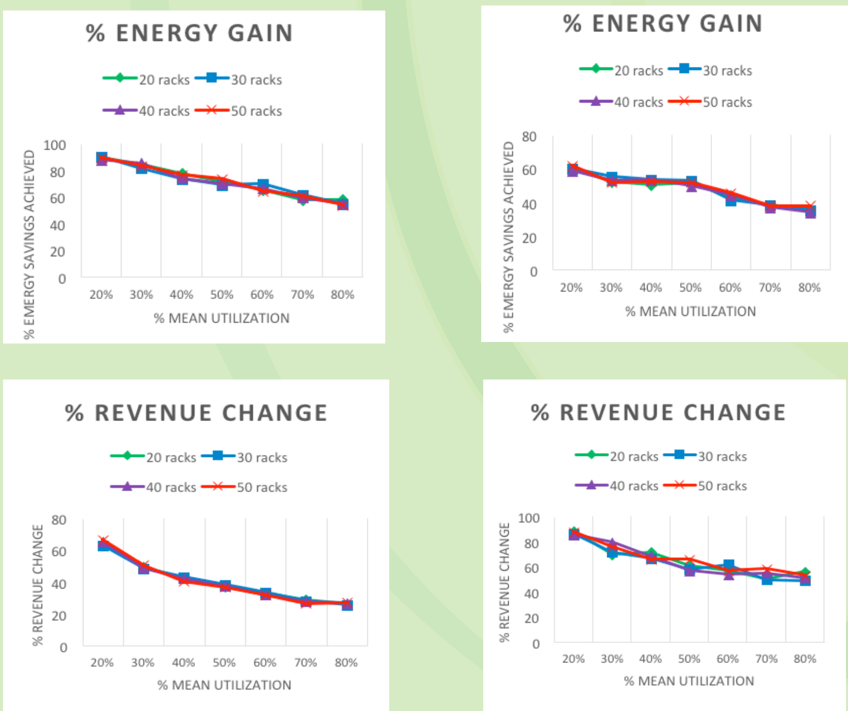

Figure 10:DOLFIN performance on medium-size urban DCs when the DC policy is set to optimize against energy efficiency (left column) or performance (right column).

of medium-sized DCs operating against 20 – 50 racks, hosting 100 – 250 physical servers and 10% of the DC infrastructure powered by green sources, **the performance of DOLFIN is able to provide energy-efficiency and revenue increase of over 80% in certain cases (very low utilization).** Even in the case of highly-utilized DCs, the energy and revenue benefits acquired through DOLFIN are very significant, reaching 60% in the case energy or performance policy is applied, respectively. In general, the effect of DOLFIN in partially green-powered medium-sized DCs is shown to be greater when compared to small ones.

4.Strengths and Opportunities for stakeholders

DOLFIN, as a whole, can be viewed as an enabler for achieving better integration of the urban DCs to their surround Smart Cities by means of optimal energy efficiency management and resulting in providing to the DCs the capability of acting like active Smart Grid actors offering balancing services to the Smart Grid operators, while maintaining their customer satisfaction at high levels through employing flexible SLAs.

Strengths

Innovative (unique) and holistic approach
The main distinctive asset of DOLFIN is the novel approach proposing to increase and improve the collaboration between synergetic DCs, energy providers' smart grids, and DCs' customers, with the common goal of saving energy and reducing carbon emissions. In addition to that, new concepts, like the Green SLA, are being developed during the three years of the project. The innovation brought by DOLFIN makes the solution unique with respect to any other present on the market.
The DOLFIN solution could increment the Net Present Value (NPV), according to the energy cost saving model adopted: in fact, compelling benefits exist, such as more efficient DC management, CO2 emissions reduction, better informed capacity planning. Although these cannot be monetized as directly as energy reduction, they make the evaluations incrementally attractive and this is another important result of this innovative approach.

No installation investment
The DOLFIN solution does not require any new or costly hardware equipment as it can run on a common VM placed on the customers' platform, therefore the initial investment of installing the system is zero

CAPEX deferral
Using the DOLFIN solution, servers already present in the DC will be used in a more energy-efficient way, leading to savings in terms of energy consumed and thus possibly postponing the need to invest in brand-new energy efficient hardware.

Economic benefits for all users of the system: DC owners, energy providers and end users
Every user of the system will enjoy economic benefits: thanks to the DOLFIN model of active collaboration, the DCs can reduce the energy costs and the impact of demand peaks and sudden fluctuations, whilst the energy providers, through the interface between the DCs and the Smart Grids, can have more flexibility in the planning.
The end users, even if not targeted by the deployment of technology components, will nevertheless benefit by discounted costs guaranteed by the acceptance of Green SLAs.

Marketing opportunities in the eco-system ("Eco-friendly DC")
Helping customers to become a Green DC and hereby reducing their environmental impact and additionally reducing their energy costs are a marketing opportunity.

Environmental benefit (reduced CO2 emissions)
The approaches of DOLFIN lead to a reduced power consumption and thus to reduced CO2 emissions.

Portability of the solution
The DOLFIN system is agnostic of the underlying cloud or virtualization infrastructure where it runs. The system is developed in a modular way, and whenever the will or need to add compatibility to a new cloud or virtualization provider arises, the new developments needed to achieve the compatibility are limited to the DC brokers only, leaving all the core components untouched.

Able to exploit energy price variance to maximize DCs benefits
The DOLFIN system takes energy price variance as an input for Green SLAs, thus pushing clients wishing to save money to reduce their computational loads during peak hours of energy demand. This, in turn, helps maximizing the energy saving during such peak times, increasing the amount of money saved.

Opportunities

Very limited number of similar services
A very limited number of similar services are currently offered to DCs and energy providers by ICT providers in the EU market.

Promote the adoption of the service
The pilot organisations can promote the service towards their networks and the other departments.
Help to reach goals set by European legislation.
The need to improve energy efficiency in DCs is not simply a business strategy, but it is one of the most important objectives of the EU Directive 20-20-20 to combat global warming. The EU Renewable Energy Directive (2009/28/EC) has in fact agreed to reduce greenhouse gas emissions by at least 20%, to reach the 20% share of renewable sources in final energy consumption and to increase energy efficiency by 20% by 2020.

Organisational benefits
The DOLFIN Consortium can generate an increase of synergies among Energy Providers (EPs), DC owners and customers.

Better reputation
The reputation of organizations adopting DOLFIN can be significantly improved, for being green and innovative. This is difficult to quantify and depends on the specific market, social contest, national regulations, ability of the company to disseminate it, etc.

Small DCs on the rise
In the recent years the number of small DCs has been on the rise [7]. This fact increases the number of potential users of the DOLFIN system, since small DCs are the primary target of the system.

References

[1] DOLFIN, “Deliverable D2.1: Business scenarios and use case analysis,” 2014.
[2] DOLFIN, “Deliverable D2.2: DOLFIN requirements and system architecture,” 2014.
[3] “DOLFIN, “Deliverable D3.1: Data Centre energy consumption optimization platform (eCOP) Design””.
[4] OpenStack Foundation, “Home page,” [Online]. Available: https://www.openstack.org/.
[5] L. Mamatas, S. Clayman and A. Galis, “A service-aware virtualized software-defined infrastructure,” IEEE Communications Magazine, vol. 53, no. 4, pp. 166-174, 2015.
[6] “DOLFIN, “Deliverable D5.2: DOLFIN system integration & evaluation””.
[7] R. Bergman, “7 Emerging Data Center Trends,” [Online]. Available: https://www.digitalrealty.com/blog/7-emerging-data-center-trends/.

DOLFIN at a glance

-

DOLFIN is a FP7 project co-funded by the European Commission under the ICT theme (Strategic objective FP7-SMARTCITIES-2013, ICT-2013.6.2), Grant Agreement Number: 609140

Coordinator
Matteo Biancani
Interoute

WEB: *http://www.dolfin-fp7.eu/*
Email: *dolfin-info@lists.grnet.gr*