

A Generic Architecture For Demand Response: The ALL4Green Approach

Robert Basmadjian Gergo Lovasz
 University of Passau
 Innstrasse 43, Passau, Germany
 Email: firstname.lastname@uni-passau.de

Michael Beck Hermann De Meer
 University of Passau
 Innstrasse 43, Passau, Germany
 Email: firstname.lastname@uni-passau.de

Xavier Hesselbach-Serra Juan Felipe Botero
 Universitat Politcnica de Catalunya
 Jordi Girona Street, 1 and 3, Barcelona, Spain
 Email: {xavier,hjfbotero}@entel.upc.edu

Sonja Klingert
 University of Mannheim
 A5, B101, Mannheim, Germany
 Email: klingert@informatik.uni-mannheim.de

María Pérez Ortega
 GFI
 Heverlee 3001, Belgium
 Email: mportega@gfi.es

Juan Carlos Lopez
 GFI
 28033 Madrid, Spain
 Email: jclegea@gfi.es

Andries Stam Rick van Krevelen
 Almende B.V.
 Westerstraat 50, Rotterdam, Netherlands
 Email: {andries,rick}@almende.org

Marco Di Girolamo
 HP Italy Innovation Center
 Milano, Italy
 Email: marco.digirolamo@hp.com

Abstract—Demand Response is a mechanism used in power grids to manage customers' power consumption during critical situations (e.g. power shortage). Data centres are good candidates to participate in Demand Response programs due to their high energy use. In this paper, we present a generic architecture to enable Demand Response between Energy Provider and Data Centres realised in All4Green. To this end, we show our three-level concept and then illustrate the building blocks of All4Green's architectural design. Furthermore, we introduce the novel aspects of GreenSDA and GreenSLA for Energy Provider–Data centre sub-ecosystem as well as Data centre–IT Client sub-ecosystem respectively. In order to further reduce energy consumption and CO₂ emission, the notion of data centre federation is introduced: savings can be expected if data centres start to *collaborate* by exchanging workload. Also, we specify the technological solutions necessary to implement our proposed architectural approach. Finally, we present preliminary proof-of-concept experiments, conducted both on traditional and cloud computing data centres, which show relatively encouraging results.

I. OVERVIEW

With the energy consumption of ICT mushrooming for some decades, and data centres at the heart of this development, a lot of research has been dedicated to this huge problem for environmental health and resource depletion. However, it turns out that data centres are not only part of the problem but also one key to its solution because the energy challenge is both, a problem of energy consumption and a problem of power consumption: In times of low supply and high demand, extra power needs to be provided at high environmental cost, in times of high supply (e.g. through wind and sun) and low demand, superfluous energy suppliers are cut off the electricity net.

The project *All4Green*¹ shows that data centres with their huge power hunger can play a role in solving this challenge.

¹<http://www.all4green-project.eu/>

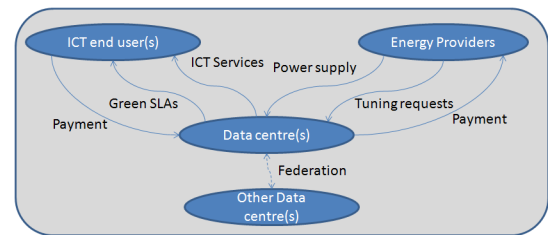


Fig. 1. An overview of the All4Green relevant ecosystem actors

To this end, the data centre is viewed as part of an ecosystem consisting of ICT users deploying services in the data centre, electrical power providers, and data centres cooperating in a federated way. By establishing a collaborative scheme within this eco-system through green contracts supported by an underlying signalling technology, All4Green tackles both goals: It aims at saving CO₂ emissions by enabling a cleaner energy mix for the energy consumption of a data centre. And additionally it will reduce this energy consumption by 10%.

All4Green relevant actors in the system, as illustrated in Fig. 1, are the following:

- Energy Providers (EP) represent the providers of energy for the data centres.
- Information Technology Customers (ITC) in the context of All4Green are not meant to be single users, but companies that make contracts for (and then potentially consume) services of the data centre.
- Data Centres (DC) act as providers of computing services for ITCs and demand energy from EPs. Data centres can federate and collaborate to exchange load among them.

The *All4Green* approach is based on a three-levels-concept

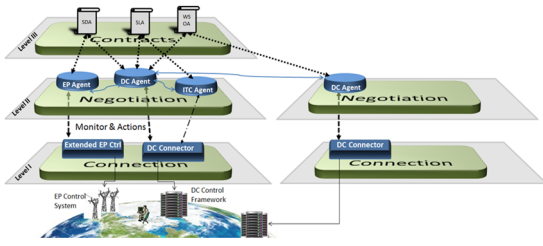


Fig. 2. All4Green design overview

with a connection level, that links the eco-system entities with the All4Green framework, a negotiation level, that negotiates the optimum solution, and a contract level containing the green contracts. In order to further elaborate, Fig. 2 depicts the All4Green design enabling Demand Response mechanism to take place between EPs – DCs – ITCs. More precisely, **Contract** level contains the contracts that organise the collaboration within the eco-system. Four different types of contract have been developed in All4Green:

- GreenSLA (green service level agreements) contracts are agreements between data centres and IT customers, which reflect the agreed scope for the data centre to operate in an energy aware way and the same time guarantee a certain level of quality of services (QoS) for the IT customer.
- GreenSDA (green supply demand agreements) contracts are agreements between energy providers and data centres, which define the flexibilities and energy-related constraints that these parties grant each other.
- GreenWSOA (workload services outsourcing agreements) contracts are agreements among federated data centres that set rules for the geographical shifting of workload.
- DC-Wide Contracts are agreements not related with a service. This agreement represents the different actions that a data centre can perform in order to reduce/increase energy. For instance, data centre’s Air Conditioner (AC) heating up/cooling down and Uninterruptible Power Supplies (UPS) discharging/charging.

On the other hand, **Negotiation** level is the kernel of the system, where the negotiation logic runs in components (called agents) that act on behalf of the main actors of the system. Agents read the contracts and monitored data from the connectors (see below for definition) and find the optimal power aware actions which are then sent to the other partners. There is a different agent for each actor in the system:

- EP agents
- DC agents
- ITC agents

Finally, the bottom level **Connection** contains all the connectors which implement the logic required for the agents to communicate with the Energy Providers and Data Centres. The connectors are the elements which connect the management framework of data centres and EP systems with the agent platform. They monitor the values provided by the management

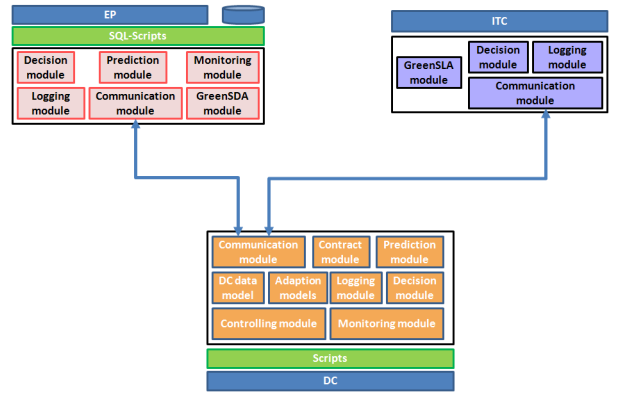


Fig. 3. An overview of the All4Green architecture

frameworks retrieving all the relevant information required by the agents, they communicate requests from different actors to agents platform, and vice versa proposals from the agents to the actors. Consequently, there is a need for a specific connector to each participating entity of type EP or DC.

This concept leads to a decentralised architecture for the All4Green system which is made up of the following building blocks (see Fig. 3):

Each actor in the ecosystem has its own sub-system which contains at least a module for the respective green contract, a prediction and a decision module as well as a logging component. Additionally all have a communication module, however, the data centre is at the centre of communication, the IT customer and the energy provider do not need to exchange information directly.

In the rest of this paper, we present the most relevant components of the All4Green architecture. We first detail in Section II the components for the Energy Provider – Data centre sub-ecosystem, then describe those for the Data centre – ICT sub-ecosystem in Section III. The aspect of federating Data centres is covered in Section IV. In Section V, we present the technological means to implement our All4Green approach. The preliminary results and conclusion are presented in Section VI.

II. ENERGY PROVIDER – DATACENTER SUB-ECOSYSTEM

This sub-ecosystem is responsible for managing the power adaption collaboration between Energy Provider (EP) and Data centres (DC). To foster this collaboration, there is a need for both common (e.g. Communication, GreenSDA) and specific (e.g. Decision of EP) modules. In this section, due to space considerations, we present only the most relevant modules of the architecture.

A. GreenSDA Module

The major inconvenience of today’s energy tariff between EP and DC is its lack of flexibility. In other terms, both parties sign an agreement that is viable on a yearly basis. Consequently, with the current energy tariffs it is not possible to enable power adaption collaboration and hence apply Demand Response (DR) mechanism between EP and DCs which requires highly flexible participants. For this purpose,

in All4Green we propose the so-called Green Supply Demand Agreement. In short, a GreenSDA specifies contractual terms that enable the power adaption collaboration between EP and DCs. For instance, one of the key contractual terms specify the amount of minimum and maximum duration as well as power the DC can increase (e.g. in case of energy surplus) or decrease (e.g. in case of energy shortage) throughout the year. More precisely, there is a contractual term that specifies for minimum power increase (decrease), its minimum and maximum durations. Similarly, there is a contractual term that specifies for maximum power increase (decrease), its minimum and maximum durations. Furthermore, the GreenSDA allows the DCs to reject power adaption requests of EP by specifying the maximum number of successive rejects as well as total rejects per month. On the other hand, in order to guarantee that the EP will not abuse the willingness of DC to collaborate, the GreenSDA defines a contractual term that specifies the maximum number of power adaption requests EP can send to DC. Moreover, it specifies also the minimum duration of time between two consecutive power adaption requests. In this way, the proposed approach ensures that the DC has enough time to recover before committing to a new power adaption request. Finally, in order to limit the time during which a DC needs to send back a reply to EP, the GreenSDA specifies the maximum reaction time.

In order to create incentives for the different parties to join All4Green's DR mechanism, reward and penalty schemes are defined. The former applies to DCs that are willing to collaborate; consequently the more DC is willing to collaborate, the higher its reward is. On the other hand, any time that one of the parties breaches any contractual term, the corresponding party needs to pay penalty in addition to depleting the received reward.

B. Communication Module

We described in the previous section that GreenSDA defines the flexibilities of the participating entities to All4Green DR mechanism. However, such a power adaption collaboration is not possible without having a communication framework enabling both parties (e.g EP and DC) to exchange information between each other. For this purpose, in All4Green, we define two types of communication messages: Monitoring and adaption. The former is used in order to obtain certain information on the state of the participating parties, whereas the latter is specified for power adaption request purposes. More precisely, the monitoring messages are as follow:

- 1) DC can ask EP for its own power consumption by specifying the exact timestamp. As a reply, the EP sends back the power consumption (Watt) of the specified time period to DC. Note that such a monitoring message is necessary in case the DC has no access to the smart meter that measures its own power consumption.
- 2) EP can ask DC for the expected power consumption by specifying the duration (i.e. start time and end time) as well as resolution². As a reply, the DC sends back to EP a list of expected power values for

²Defines the time between two power consumption values in the power profile.

the specified duration. Note that DC can estimate its power consumption by means of a Prediction module (see DC component in Fig. 3).

- 3) Similar to the previous monitoring message, the DC can ask EP for the expected load curve of predicted power surplus and shortage of the grid for a given time period.

On the other hand, the adaption messages are given by:

- 1) Whenever there is a power shortage or surplus situation, the EP sends a power adaption request message to DCs by specifying the duration (i.e. start time and end time) as well as the type of the adaption (i.e. increase/decrease). In return, the DC sends back to EP at least one power profile where each such profile contains the following information:
 - The amount of power that the DC can adapt (e.g. increase or decrease).
 - The duration of the power adaption.
 - The expected additional power consumption during the recovery phase. Also, the profile specifies the time when the recovery starts as well as its duration.
- 2) After receiving a power adaption request, the DC sends back either a set of power profiles (see previous message type), or a negative acknowledgment where in that case the DC informs EP of its unwillingness to collaborate for power adaption request.

C. Decision Module of EP

In All4Green, every time that there exists a power shortage or surplus situation, the EP sends a power adaption request to DCs asking for their collaboration. Furthermore, in All4Green we adopted "request everyone" approach meaning that the EP broadcasts all the participating DCs for collaboration. In return, the EP will receive from DCs a set of power profiles (at least one) or a negative acknowledgment. However, in order to choose the most appropriate profiles of DCs, policies are needed. A policy defines strict rules of how to select profiles of DCs. In All4Green, we propose a "fairness" policy which has the objective of distributing evenly the burden of power adaption collaboration among the participating DCs. To this end, the Decision module of All4Green's architecture is devised in order to implement such type of policy and even have the possibility to add other policies in the future.

D. Adaption Models of DC

In order to be able to decide which action or combination of actions can be taken when the EP sends an adaption request, the data centre has to be able to decide how these actions would affect the overall state of the DC, before actually applying them. The DC can simulate how various adaption strategies would influence the operating conditions and whether they would meet all the safety requirements of the DC. Therefore, for each adaption capability, the DC has to refer to specific models to simulate the results of several adaption actions. Since each model is highly specific depending on the modelled adaption capability, we shortly list the main properties of each approach:

- UPS

Since the UPS is a fallback device which the DC relies on during power *blackout*, the DC usually does not want to use its full capacity during a power *shortage*. Therefore, the model has to predict the maximum/minimum time the DC can run the UPS until the UPS has lost more than a pre-defined amount of power. Battery capacity also depends on several criteria, e.g.,

 - o average room temperature
 - o age of the battery
 - o the number of discharge cycles
 - o the average depth of discharge
- Air conditioning
 - o architecture of the DC building e.g., whether servers are hosted in an hot aisle/cold aisle environment
 - o size of the room(s)
 - o speed of the server fans which influences air volume that is moved in the room(s)
 - o maximum/minimum temperature boundaries of the room(s)
- QoS of ICT services

The QoS models have to map the impact of QoS degradations to the power consumption of the servers.

Necessary, common features of all models are:

- Computation of the impact:

The models compute the impact of adaption capabilities. E.g., the minimum/maximum duration the DC can use a capability in order to reduce its power consumption can be derived, and the amount of power that can be saved.
- Computation of recovery duration and recovery power:

Models compute how much time and power will be needed to go back to the original state of the DC after a specific adaption capability has been performed. E.g., the power and duration that is needed for recharging the UPS after it has been recharged to a certain amount.
- Compliance of security requirements:

Models ensure that all safety requirements of the DC are covered and that none of them is violated. For example, the model of the air conditioner ensures that the maximum, pre-defined temperature of the DC is never exceeded.
- Compliance of contractually agreed requirements:

Furthermore, models have to cover all demands that were contractually agreed on. E.g., QoS of some services should never be reduced by more than a specific degree.

III. DATACENTER – IT CLIENT SUB-ECOSYSTEM

This sub-ecosystem refers to the elements and functions to permit the collaboration between the ITC and DC in order to reduce the power consumption and CO₂ emissions by

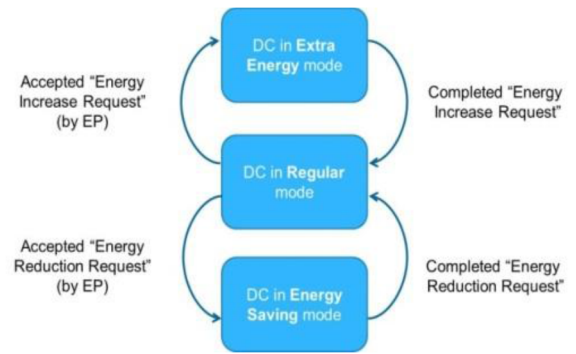


Fig. 4. DC Energy modes context in All4Green

energy efficiency strategies within the DC, and to support the EP in avoiding power shortages, integrating renewable energy sources by enabling a collaboration between EP and DC. In this section, the essential modules providing those functionalities are presented.

A. GreenSLA Module

The liabilities between a DC and its customers are ruled by a set of contracts. Apart from framework contracts, for the delivery of each DC service, a Service Level Agreement (SLA) is closed. The so-called GreenSLA extends the existing SLA agreements in order to include energy and carbon emission aware parameters. Therefore, in All4Green, GreenSLAs are extensions to traditional SLAs, including three main additions: *Flexibility*, *GreenKPIs* and *Collaboration*. The flexibility is the variability that the ITC and the DC are willing to accept in each of the service's running conditions requirements (e.g. performance, availability, execution, maintenance).

Let's define a *context* as the specific situation or state that allows the possibility of modifying the service conditions to promote a more environmentally friendly behaviour. The *flexibility* considers changes in context in order to take decisions.

In this project, 3 contexts are considered: Time/Calendar, DC Energy Mode, and Load.

Time / Calendar context is related to a time period: hour(s) and/or day(s) of the week and/or the year's season. An example of a time-dependent GreenSLA clause could be: High Availability + High Performance in working days, and low availability + low performance during nights and weekends.

DC Energy Mode context is related to the current energy mode of the DC. Generally, the DC is in Regular Energy mode (RE), but it can change to Extra Energy (EE) mode or to Energy Saving (ES) mode for a certain time period due to the reception and acceptance of a power adaption request coming from the EP. Other modes can be considered in the future, such as Emergency Mode, in order to include specific situations.

Load context is related to the behaviour of the service-load during different time frames, in order to change the configuration of the services depending on the real-time load. This change of configuration should be performed trying to promote energy or CO₂ emission saving. The concept can be exploited by load balancers to consolidate services in a reduced amount of servers, to shut down some of them.

The performance measurements under the concept of the *GreenKPIs* are new service level indicators in order to evaluate the success of the goals for energy consumption and/or CO₂ emission reduction. For instance, a GreenSLA could guarantee a certain CO₂ emissions level during the execution of an IT service by allowing continuous interaction between the DC and the EP.

Collaboration in the DC–ITC sub-ecosystem is regulated by the GreenSLA and refers to the requests exchanged by the ITC and DC. These requests coming from the DC are triggered by the change of the DC mode: From regular mode the DC might switch to situations where it needs to reduce its current energy consumption (energy saving mode) or, on the contrary, switch to a situation where it needs to consume extra available energy (extra energy mode). This change of DC status comes from the interaction of the actors in the DC–EP sub-ecosystem. Also, from ITC to DC, the ITC can share its plans of demand and, taking that into account, the DC suggests to the ITC how to organise this in a way that the demands are fulfilled and the energy consumption is reduced. For this purpose, the GreenSLA may contain clauses ensuring that the ITC shares its plans, something like: Every timeslot the ITC will inform about the expected behaviour (load demand) of the service during the next timeslot. From this, the ITC will have a monetary bonus, and, on the contrary, if it does not provide the expected behaviour, the user will be charged with a monetary penalty.

B. Communication Module

The mechanisms and strategies provided by the GreenSLA requires a communication framework to exchange the information in the DC–ITC sub-ecosystem. Interfaces are defined between an agent for the DC and an agent for the ITC. As in the EP–DC sub-ecosystem, the messages can be: Monitoring (to get information and to estimate consumption of services or servers) and adaption (to perform power adaption requests between the DC and the ITC), but also others are included in the DC–ITC sub-ecosystem: Management (services information, warning, failures or alarms) or collaboration (to perform collaboration requests).

C. Decision Module of IT

In order to take decisions concerning services execution in the DC, the corresponding agent (for each data centre) is able to get the necessary information to provide a Service Table, to summarise the description of the running services. From this table, a set of policies defining rules will implement the intelligence. The evaluation of a policy will provide a list of actions to be executed regarding the configuration of IT Services which may ultimately translate into lower-level specific actions to be executed by the DC through a specific interface with the aim of fulfilling the aforementioned general goal.

IV. FEDERATION OF DATACENTERS

A considerable impact on both energy consumption and CO₂ emission savings can be expected if data centres start to *collaborate* by exchanging workload. Typical examples of these can already be found in practice in existing *single*

owner data centres: multiple data centres owned by a single entity, which distribute and exchange workload according to user needs or in order to optimise resource utilisation. Within All4Green, such federations of single owner, but also *multi owner* data centres are studied, whereby workload allocation, outsourcing and insourcing takes place according to energy and emission saving demands. To this end, we have developed a *Workload Services Outsourcing Agreement Module* (WSOA Module) for the establishment of data centre federations, and several *negotiation policies* for outsourcing and insourcing workload.

A. GreenWSOA Module

A GreenWSOA is an agreement between two data centres that intend to collaborate in improving each other's (green) performance/efficiency by exchanging workload. By committing to a GreenWSOA, the collaborating data centres thus become a *federation*. A GreenWSOA can be considered as a unilateral "advertisement" of one DC informing another remote DC what kind of collaboration it can offer. GreenWSOA agreements are generated according to local templates and entail specification of a federate data centre's *capabilities*, *compliance*, and *service compatibilities* with respect to another data centre. Capability terms include for instance "remote allocation" and "live migration", while compliance terms may concern specific security or privacy certification. Service compatibilities map local service names and options to the names used at the remote data centre. By means of this lightweight approach, starting with unilateral agreements between data centres, larger federations can gradually come to existence without forcing individual data centres to adhere to any externally imposed standard regarding services and service options. Furthermore, at all times data centres control the set of data centres with which they intend to collaborate.

Based on the terms specified in the GreenWSOA contracts agreed upon by federated data centres, their representative DC Agents can enter into negotiations about outsourcing or insourcing some of their workload. Such negotiations may occur either (i) when some data centre receives a request from its energy provider to increase or decrease its power consumption for some interval in the future, or (ii) when the data centre wishes to improve its local performance based on expected work load, or alternatively (iii) based upon a coordinated incentive to save CO₂ emissions in a situation where the insourcing data centre can use green energy while the outsourcing data centre cannot. Furthermore, depending on terms agreed with customers in the respective GreenSLAs, workload may be either outsourced during initial allocation, or possibly "live" migrated during execution.

B. Negotiation Policies

Within All4Green, two approaches to negotiation within the data centre federations are considered. The first is for the initiating data centre to contact all known federates' agents directly and request quotes or offers for allocating or migrating some services over some interval. The second is to introduce a *broker agent* that maintains some abstract status information on all federates regarding their current availability or "willingness" to accept new workload or to outsource current workload. The main advantage of introducing a broker

providing likely candidates for insourcing and outsourcing is the increase in negotiation responsiveness and reliability. This becomes particularly important in the case of larger federations or critical scenarios (e.g. sudden peak loads in the local power grid) where speed is of the essence. Furthermore, a broker is able to act as a trusted third party, thereby removing any necessity for data centres to share possibly business sensitive information with other (competitor) data centres.

Either way, the data centre initiating the negotiation compares the available options and selects the best one, where “best” is a relative term that depends on the local data centre’s preferences. One possibility is to apply multi-criteria analysis where the criteria considered may include (i) the amount of emissions, (ii) the economic price, and (iii) the power consumption costs involved with the reallocation or migration of workload for each candidate data centre in the federation.

V. IMPLEMENTATION

GreenSDA, GreenSLA, GreenWSOA and DC-Wide contracts are implemented using and extending Webservice-Agreement for Java (WSAG4J) framework. WSAG4J framework allows to represent agreements using XML format and provides automatic monitoring of the contractual terms of the green agreements, raising events when the terms are violated. WSAG4J framework stores templates for each type of contract that are instantiated as concrete agreements.

Almende’s Eve Agents technology is used to implement the decisions support system as well as the communication framework of All4Green software. There are two types of agents considered in All4Green: *Delegate* and *Wrapper* agents. The *Delegate* agents represent the energy providers (EP Agents), data centres (DC Agents) and IT customers (ITC Agents). They contain the implementation of the decision support system (e.g. logic) for behaviour on their behalf. Delegate agents communicate with **Contracts** level (see Fig. 2) using *Rest* to retrieve the information inside the instances of the agreements contained in WSAG4J framework. Connector *Wrapper* agents manage the energy provider and data centre connection so other agent platform levels can interact with them in a uniform way. Agents use JSON/RPC to communicate with each other including communication between Delegate and Wrapper agents.

The communication with the specific infrastructure of Data Centres and Energy Providers is done with pieces of software called *Connectors*. Connectors are tailor-made for each Data Centre or Energy Provider. Connectors monitor information from the DC or EP using different technologies depending on the specifics of DC or EP infrastructure. For instance, as an example SNMP, SSH, WEB services are technologies currently used inside connectors. The Wrapper agents, that are part of the connectors, provide the monitored information when requested by Delegate agents and communicate the suggested actions, sent by Delegate agents, to the connectors that interpret and send the actions to DC or EP when needed.

VI. PRELIMINARY RESULTS

The All4Green concept will be evaluated and validated by two trial cycles, whose goal is to assess at what extent the All4Green prototype is meeting its expected results. The

trials will occur in three different testbeds, specifically selected to cover the whole span of scenarios and specific functions implemented by All4Green.

In general, the targets set by All4Green can be recapped as follows:

- 1) For an individual data centre, a 10% reduction of energy consumption on top of any pre-existing strategies and policies pursuing the same goal (e.g., the FIT4Green³ plug-in), exploiting the establishment of GreenSLAs with the data centre users. This energy saving target increases up to 20% during certain time periods in particular conditions.
- 2) For an energy provider, enable a temporary 10%–20% reduction of the data centre energy demand, thanks to the establishment of GreenSDAs with data centres supplied by the provider itself.
- 3) For two or more federated data centres, allow to avoid unfavourable energy conditions by moving workload to the site in the best current situation, achieving reductions in the total energy consumption and/or in the total GHG emissions.

The testbeds selected to evaluate the above scenarios are the following three:

A. Traditional Data centre – Energy Provider testbed

This testbed puts together a data centre delivering traditional IT services with its energy provider. It allows to test the Scenarios 1 and 2. It is realised in the German town of Passau, at the premises of: a:k:t: Informationssysteme AG⁴ as data centre, and Stadtwerke Passau⁵ as energy provider. The Scenario 2 will be evaluated here in its different aspects, assessing the effect of collaboration in peak load detection and management, metering on the different layers of data centre operation, connection management between data centre and energy provider.

B. Data centre Federation

This testbed is devoted to assess the Scenario 3, understanding what advantages can be taken by the availability of two federated data centres different by geography and energy suppliers. It is realised at the premises of the Italian telecommunication operator Wind, whose two traditional-type data centres are located in the cities of Rome and Ivrea. Scenario 3 will be evaluated here in terms of achievable consumption decrease, and effect on emission reduction depending upon the emission factor delta of the different energy providers.

C. Cloud computing Data centre

This testbed will assess all of the three scenarios in a computational environment different from the one in the previous two testbeds. Cloud computing has different characteristics compared to enterprise IT, in terms of load profile variance and predictability, which makes of interest to evaluate the scenarios in such an environment. It is realised at the premises

³<http://www.fit4green.eu/>

⁴<http://www.akt-infosys.de/en/unternehmen/wir-ueber-uns.html>

⁵<http://www.swp-passau.de/>

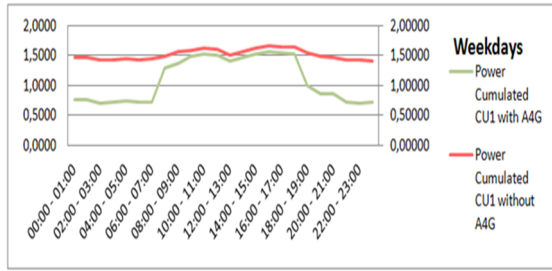


Fig. 5. Results of traditional testbed

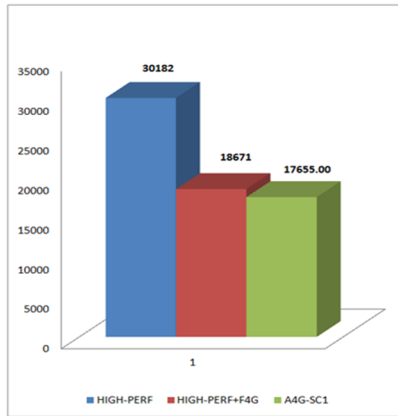


Fig. 6. Results of cloud computing testbed for Scenario 1

of HP Italy Innovation Center, with a lab-grade infrastructure including two federated data centres. Some specific analysis will be performed in this testbed.

In parallel with the development of the prototype described in Section V, a set of “manual mode” tests have been executed in the three testbeds in order to pre-validate the All4Green concept. Partial customised and quick implementations of the All4Green functions have been enacted in each of the testbeds, to evaluate the potential effect of applying All4Green’s principles to the same platforms where the the prototype is being trialled. These tests had an overall positive result, proving that the All4Green concept has the potential to meet the set numerical targets. Next, we present a quick resume of the results for this test cycle.

D. Obtained Results

Scenario 1 tests in both traditional and cloud testbeds showed very promising figures, and demonstrated that the application of GreenSLAs can induce additional energy savings on top of existing strategies. In Fig. 5, we can observe the marginal savings observed in the traditional testbed during weekdays. The employed test cases covered a range of three different service types, and in each case the results were positive. Measured savings float between 4% and 13%, and during weekends (when conditions are more favourable) savings up to 50% could be measured.

Fig. 6 shows the results of Scenario 1 evaluation in the cloud testbed. Expected results here are lower, since there is not the same level of predictability to exploit and more safety

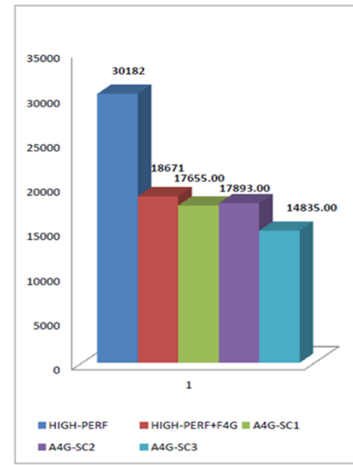


Fig. 7. Results of cloud computing testbed for Scenario 2

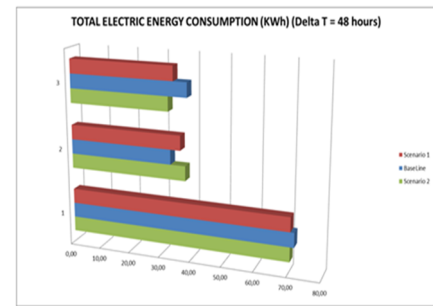


Fig. 8. Results of cloud computing testbed for Scenario 3

buffering must in general be applied. We observed an average marginal saving around 5%, rising up to 40% in specially favourable conditions.

As far as Scenario 2 is concerned, this pre-evaluation was done in the cloud testbed, and results are visible in the Fig. 8. Here, the main focus on assessing the entity of what we call the recovery effect. In short, to temporarily decrease the consumption for helping to squeeze demand peaks, the used action is a time shift of workload execution. This implies that the workload can be executed later, at a time when energy consumption conditions are less favourable, and overall a price in increased consumption is paid for being able to contrast the peaks. The measurements showed an average value of this increase around 1.3%, meaning that the Green SDA application is not bringing any significant overhead to the energy consumption status.

Finally, Scenario 3 was evaluated in the federation testbed. This pre-evaluation assessment was quite simple, just measuring the effects of rebalancing the load between two federated data centres starting from a fully unbalanced distribution. The tests showed an extra saving around 1.7% obtainable by rebalancing the load distribution inside the federation, and emission savings between 2.5% and 3% (by simulating different emission factors).

VII. CONCLUSION

Demand Response (DR) mechanism is becoming more and more prevalent for energy providers especially with the advent of the renewable energy sources (e.g. photovoltaic) which have intermittent behaviour. Data centres, due to their significant energy use, are excellent candidates to participate in DR programs. In this paper, we introduced a generic architecture that enables DR mechanism to take place between energy provider and data centres. To this end, we presented the necessary building blocks enabling this power adaption collaboration such as GreenSDA, GreenSLA, GreenWSOA contracts. We also presented the technological choices of implementing the presented architecture for DR purposes. Finally, as a proof-of-concept, we conducted several field tests both on traditional and cloud computing data centres. Although, the presented results are preliminary, however they prove the correctness of the initial assumptions and provide insights of the operational characteristics of data centres.

ACKNOWLEDGMENT

The research leading to these results was supported by the European Community's 7th Framework Programme in the context of the ALL4Green project.