2015 3rd International Conference on Future Internet of Things and Cloud

Using green KPIs for Large IT Infrastructures' Energy and Cost Optimization

Boban Celebic, Ruth Breu
Institute of Computer Science, University of Innsbruck, Austria
{*boban.celebic, ruth.breu*}*@uibk.ac.at*

***Abstract*—This research paper focuses on reengineering and optimization of energy consumption and other costs for large IT infrastructures, with support from EA software which is based on advanced use of IT Key Performance Indicators (KPIs) and powerful visualization solutions. Use of KPIs is crucial in achieving energy efficient, *green* IT; implementing these metrics allows the enterprise to determine areas to improve energy efficiency, monitor that the data center operators improve the designs and processes over time, to automate the energy and cost (optimization) management and to predict and discover opportunities to repurpose and redirect energy for additional IT equipment. Such approach allows to measure current and predict and plan future power usage, so that current IT infrastructures can be optimized or new IT infrastructures can be planned to be energy efficient from the start.**

***Keywords*-energy; consumption; optimization; energy efficiency; EA; EAM; KPI; green IT; IT metrics.**

I. Introduction

Enterprises are lately more and more forced to manage energy and resources in a sustainable and economical manner, for manifold reasons: besides legal regulations and societal impact, energy costs of IT and data centers' operations are significant, whether for internal corporate IT operations or as part of IT outsourcing. "Inefficient" equipment operations, e.g., data servers "spinning" when no active operations are being performed or processors consuming energy while doing no calculations, and similar, are among the reasons for high energy costs. In addition, equipment costs have been significantly reduced lately, causing shifting of focus on other costs - primarily energy. As a result, enterprises invest more efforts and funds in finding new and innovative ways for optimizing energy spending. This optimization implies several tasks: rethinking and reshaping the entire long-term vision and operational model of an enterprise; finding and removing weak spots in the current business model and reengineering the business processes in respect of efficiency [1]; and optimizing the (IT) infrastructure in terms of spatial, resource and energy efficiency. These tasks are complex and pose severe and variegated challenges[1], concerning long term consequences of related changes. *Enterprise Architecture* (EA) helps in solving these tasks; it is the basis for evaluating and optimizing energy usage. The purpose of EA is to achieve greater alignment between business' and IT concerns. EA achieves this goal, in part, by guiding the process of planning and designing IT infrastructure of an enterprise with the goal to meet the desired organizational objectives. These objectives, in short, consist of improving the manageability, efficiency, agility and effectiveness of the business. In order to meet these objectives, enterprise architects can also consider changes concerning other categories of the business: innovations in the structure or processes of an enterprise; the integration and/or standardization of business processes; improving the quality and timeliness of business information; and inovations in the use of Information systems (IS). Business and technology parameters can be improved through EA as well.[2]

Typically, when observed in wider context, EA encompasses all aspects of the enterprise; in more specific context, the architecture decisions are limited to the IT aspects of the enterprise - other aspects only serve as inputs. In that sense, Enterprise Architecture Management (EAM), more distinctively, clarifies how IT contributes to raising efficiency of business processes, reducing business risks, cost savings, and other [2]. We focus on this aspect; more precisely, on IT and data centers reengineering supported by use of Key Performance Indicators.

The rest of the paper is organized as follows. Section II gives a brief introduction to our approach to the problem and conducted activities. Sections III and IV outline the solution framework (green EA model), including discussion on "green" KPIs, and the support software tool. Next session discusses validation of the concepts and project. Finally, we conclude and outline future work.

II. Approach

Our project activities, conducted within the EN-ACT[3] project, have primarily focused on reengineering and optimization of energy consumption and other costs for *large IT infrastructures* and *data centers*[4], based on advanced use of green IT KPIs. The main objective is to provide enterprises

[1]Some of these challenges are related to energy usage, like: *How does improvement of data center's energy efficiency impact the whole organization? What is the long term impact on our customer service? Are the results consistent with the overall business goals?* [1]

[2]E.g., the so called ASSIMPLER critical parameters: availability, scalability, security, interoperability, maintainability, lower cost, extendibility and reliability.

[3]http://www.en-act.eu

[4]This direction is justified with the following facts:

- The ICT industry is responsible for approximately 2% of the global CO_2 emissions (IT emissions will rise from 3% of the total global CO_2 emissions in 2012 to 6% in 2020 [3]).
- Data centers alone produce 0.3-1.2% (depending on source) of world's CO_2 emissions. 1.2-2% of energy in the USA is spent by data centers, which corresponds with >0.8% of the total energy use in the world; [3]
- IT adds up to 40% of the average corporate electricity bill; IT electricity demand is ever-growing; [3]
- Simple no-cost decisions made in the design of a new data center can result in savings of 20- 50% of the electrical bill, and with systematic effort up to 90% of the electrical bill can be avoided [4];
- Server electricity costs will exceed purchase costs on useful life [3];
- The latest energy efficient servers have reduced power consumption and improved performance; Upgrade of a server can reduce energy consumption by 15%; [3]
- The industry average for server utilisation is only about 20%, with 32% running at or below 3% peak utilization. Through the implementation of the latest virtualization techniques the utilisation factor can rise up to 80%. Enterprises can reap up to 65% reduction in server count through virtualization. [3]

978-1-4673-8103-1/15 $31.00 © 2015 IEEE
DOI 10.1109/FiCloud.2015.86

IEEE computer society

with a specialized energy management software for planning and optimizing IT infrastructures based on configuration, (current and forecasted) energy consumption and other costs[5]. This energy management software will be based on *txture*[6], an EA tool for textual IT-Architecture documentation and analysis. Within the project we have developed concepts for this prototypical module (in form of a plugin for the txture tool), as well as a prototype of the interface. The energy management module should support the specification, forecast, optimization and monitoring of energy-related goals and metrics in an IT infrastructure environment. In such a way, the aspect of energy consciousness in the management of IT-related metrics comes into focus. Optimization of IT infrastructures, with the help of the energy management module, will be based on proposed green EA model (described in the following section) and (identified and categorized) green IT related KPIs, supported by powerful and intuitive visualizations.

In order to achieve our main goal, the following activities were conducted:

A. Enterprise Architecture reengineering and EA model

First, existing business models were investigated in order to find and remove the weak spots in those models and to reengineer the business processes in respect of efficiency. The objectives consist of improving the manageability, efficiency, agility and effectiveness of IT applications in the context of the business. In the continuation, we have focused on IT and data centers reengineering supported by use of KPIs.

B. Investigation and evaluation of existing IT KPIs

During our in-depth preparation and investigation in the previous phase, we have realised that the use of KPIs is crucial in achieving energy efficient, "green" IT[7]. Thus, an extensive literature review was conducted in order to collect, compare and evaluate all existing "green" IT KPIs, and to decide on which ones would be suitable for use in our approach and, eventually, tool. These KPIs play the strong role in our EA approach and model, as explained in the following sections. In addition, IT experts should have the possibility to use these metrics to monitor and optimize any aspect of the processes within the IT infrastructures.

C. Strategy for data center energy efficiency optimization

The strategy for achieving energy efficiency within data centers consists, in short, of: determining the energy baseline, forecasting IT growth, data analysis and, finally, energy management. These steps can be translated into these activities: 1. Define data center efficiency. Define energy efficiency metrics. 2. Analysis of current environment. Use of energy efficiency metrics. Guidelines for future power reduction. 3. Build & employ data center efficiency strategy. These principles will be explained in more detail in subsections III.b and IV.a.

[5]The final goal is out of EN-ACT project scope and funding.

[6]http://txture.org

[7]Implementing these metrics allow the enterprise to determine areas to improve operational efficiency, designs and processes over time, to automate the energy and cost (optimization) management and to predict and discover opportunities to re-purpose energy for additional IT equipment. Such approach allows to measure current and predict and plan future power usage, so that current IT infrastructures can be optimized or new IT infrastructures can be planned so to be energy efficient from the start.

D. EA/SOM tools review

After the EA and KPI phases, we conducted a thorough research of numerous existing EA and SOM (System and Operations Management) tools and their approaches to measuring and optimizing IT energy efficiency within enterprises and data centers, to demonstrate that there is indeed a strong need for an energy optimization software which could not only monitor and measure energy usage, but also forecast, compare, plan and optimize future energy usage within large IT infrastructures and data centers. We found out that very few of the available tools pay attention to IT energy efficiency optimization; in addition, their functionality is usually, in best case, limited to (often rudimentary) measurement and monitoring of energy consumption. We have found no tool providing capabilities for forecasting and planning new or optimizing an existing IT infrastructure in sense of energy efficiency. Thus, our research focus was directed towards the development of an energy management module for an existing EA(M) tool (txture). We are confident that this proposed energy optimization and management module amends the missing functionalities of existing tools and that it will be able to reduce total energy consumption in any IT infrastructure.

In addition, the review of existing tools had another goal - to help us formulate concepts and requirements for our own energy optimization software. In this sense, we evaluated the approaches for using KPIs in existing EA tools, with aim to develop novel ways of integrating KPIs within the IT infrastructure management perspective and energy software.

E. Concepts and Requirements for the energy management software

After the previous investigative activities, we started developing concepts for the innovative energy management module (more on these concepts in section IV).

F. User interface prototype

At the end, we have developed a graphic user interface prototype for the energy management module, based on use of advanced visualizations.

Results and outcomes of these activities are described in the following sections.

III. Green Enterprise Architecture Model

We introduce our EA (meta) model[8] by showing its different layers and their relations in more detail. These layers are well integrated and depend on each other. They show the interactions between the technical resources of the enterprise and the business level going through the software environment. Figure 1 shows the three layers (technical, information and business) of the EA meta-model. The meta model is depicted bottom up in order to ensure a holistic overview which describes the evolution of a business process from the needed resources up to the realization. All three layers are explained below.

Technical Layer is the lowest level of the meta-model. Infrastructure elements can be components such as servers, networks, storage units, devices, CPUs, etc. These components communicate with each other and enable information systems to run on them.

Information Layer consists of information systems, applications and IS objects. Concerning the direction of the

[8]An EA model should be able to address various questions on energy consumption, IT hardware usage efficiency, strategic planning and other. [1]

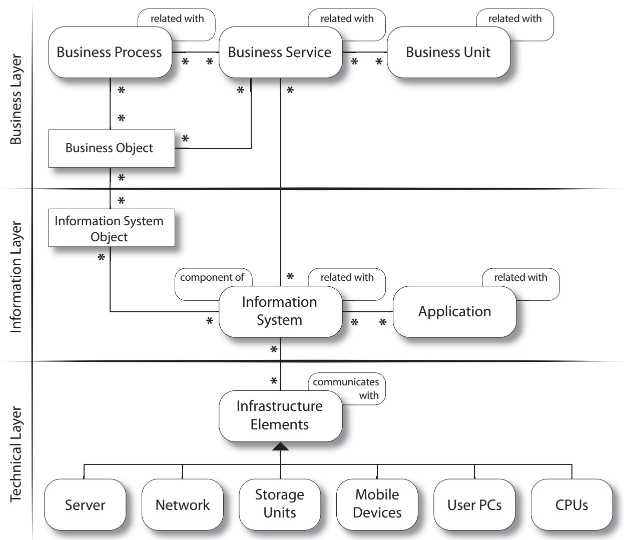

Figure 1. Layers of the "green" EA model

connection, IS objects are gathered from or converted into business objects. Information systems are connected to the entity *business service* from the business layer, as they are needed to fulfil the business services. IS objects are used as an input and are also produced as an output of IS. Concerning the architecture of the connected systems, the "*component of*" connection suggests tight coupling of systems - with one system being a subsystem, while the other system being a super-system. In addition, the connection "*related with*" indicates loosely coupled information systems. Lastly, there are applications which run on one or more information systems. Applications can also be related with each other.

Business Layer describes the relations among high level elements of the business view. There are only m:n[9] relationships among the business object, process, service and unit. A business process may be associated with many different business objects and a business object can be used by multiple business processes.

To monitor how efficient different layers are, key performance indicators were defined.

A. Key Performance Indicators

Crucial part in any energy and cost optimization strategy are the Key Performance Indicators. KPIs are metrics[10] whose role is to help monitor various parameters of importance to the optimization strategy. In order to monitor the effectiveness of the business processes, specific KPIs have been identified for each layer[11]. Table I summarizes the categorization of KPIs by layers in which they can be applied.[12]

Table I
OVERVIEW OF LAYER SPECIFIC METRICS

Technical Layer
Power Usage Effectiveness and Data Center Infrastructure Efficiency [8]–[11], Green Energy Coefficient [9], [11], Energy Reuse Factor [9], Carbon Usage Effectiveness [9], [11], Deployed Hardware Utilization Ratio [10], Deployed Hardware Utilization Efficiency [10], Compute Power Efficiency [10], CO_2 Emission [10], IOPS/Watt [10], [12], Bandwidth/Watt [10], Capacity/Watt [10]
Information Layer
Response Time [12], Energy Aware Application Performance [12], Process Time/Job Duration [10], Throughput [10], Availability Rate [10], Reliability [10], Recoverability (Repairability) [10], Application Performance Indicators [10]
Business Layer
Green Energy Coefficient [9], [11], Human Resource Indicator [10], Compliance Indicator [10], Infrastructural Costs Indicator [10], Carbon Credit [10], Return of Green Investments [10], Paper/Digital Ratio, Paper Print Indicator

B. Data centers

As mentioned, we focus on IT and data center reengineering, supported by use of KPIs. The question is how exactly can energy usage in a data center be reduced? As hinted in section II, strategy for optimizing energy consumption within a data center consists of the following steps (principles):

- First, data center efficiency needs to be defined. This is achieved by identifying and defining the initial set of energy efficiency metrics, and also implies determining the energy baseline. Once the energy baseline is determined and process for measurement of each metric is defined, it becomes possible to measure the result of any improvement while gaining the ability to pinpoint problematic areas fast;
- Second, IT growth needs to be forecast. Good forecasts lead to better planning and more effective solutions.
- Then, analysis of the current environment needs to be performed. This implies, among other, examination of infrastructure adjustments to maximize energy efficiency. This analysis should result with a list of guidelines for future power reduction;
- Last, data center efficiency strategy needs to be defined and employed. This implies constant monitoring and management of energy consumption and application of necessary measures[13] for correcting potential imbalances. As stated before, KPIs should be set for energy consumption and other costs. Use of these KPIs will lead to suggestions for improving energy efficiency.

[9] *Many-to-Many* relationship

[10] A "good" (i.e. effective, efficient) IT KPI should satisfy the following requirements [5]–[7]: it should have a clear and understandable name; must be technically sound (scientifically accurate and used precisely), easy to use (intuitive) and measure (can be easily calculated from accessible data on IT equipment), and cheap for implementation and operation; clearly shows the impact of and is related to key business objectives; provides a clear, intuitive understanding of the impacts of changes; should be comprehensive enough to evaluate the whole system, without overlooking any component, yet granular enough to analyze individual aspects as well; should be capable of scaling according to the purpose for which it was initially created and should factor technological, economic, environmental changes; can be designated as a percentage (dimensionless) with an ideal of 100%; should be capable of identifying key improvement areas and providing data driven decisions, by supporting "what if" analysis for IT planners and operators.

[11] Within the *technical layer*, KPIs mainly take hardware aspects into consideration. Within the *information layer*, Quality of Service (QoS) aspects are the most important. These KPIs are mainly concerned with quality of service and indicate the trade-off between energy efficiency and quality. Quality arguments like response time, throughput rate, reliability and recoverability are considered at this level. The *business layer* is concerned with the business processes needed to fulfil the business services for the business units. This top layer mainly affects the underlying layers and concerns aspects like return of green investments, media disruptions, optimization of processes and the choice of the energy vendor.

[12] Thorough discussion on the proposed EA model and framework, its layers and KPIs, is beyond the scope of this paper.

[13] Decommissioning or repurposing servers which are no longer in use; powering down servers when not in use; replacing inefficient servers; virtualizing or consolidating servers, and other [13].

What needs to be emphasized, in regard to previous principles, is that they all must be driven by metrics (KPIs), standards and derived guidelines.

IV. Energy & costs optimization module for the Txture tool

The concepts, results and prototype of the user interface for the energy management software should serve as a basis for future development of energy & costs optimization module for *txture*, an EA tool for textual IT-Architecture documentation and analysis with ability to cope with 100.000+ documentation items. Txture (currently in version 2), consisting of a multi-user Eclipse-based modelling tool and a web-app to flexibly visualize the architecture for different stakeholders in any organization, is already positively recognized in industry. A tool demo is available at http://txture.org. With the envisioned energy optimization module, txture should gain functionality for IT infrastructures' planning, prediction and optimization based not only on standard criteria but also on energy consumption, performance and other costs, and thus should provide enterprises of all sizes with a distinct possibility to plan their energy consumption ahead or to optimize the existing consumption and costs.

A. Energy module's functionality

As stated before, we have developed concepts and a prototype of the graphical user interface for a novel energy and costs optimization module for the txture tool. In this section we outline these concepts, use scenarios and sketch the module's functionality, which is addressing the challenges described in previous sections.

Typical *scenarios for optimization* (of energy consumption and other costs) of IT infrastructures we consider are[14]:

1) Planning of a *new* IT infrastructure based on specific requirements which have potentially been implemented in some other, existing infrastructure(s). The knowledge and experience from the previous installation can be re-used for new implementation;
2) Planning of a *new* IT infrastructure, with requirements that haven't been previously implemented in another, existing infrastructure - implying that there is no possibility for re-use of previous knowledge and experience;
3) Expanding (significantly) an *existing* IT infrastructure;
4) Optimization of an *existing* IT infrastructure or its part(s) (full or partial optimization).

The *concepts* on which energy optimization module is based can be summarized in one simple principle: Plan (forecast, analyze, build strategy) & Manage (employ strategy, measure, diagnose, correct, calculate trends).[15]

The *functionality* of the envisioned energy optimization module can be summarized as follows:

[14]While many tools for measurement and monitoring of energy are available on the market, few pay attention to the first three scenarios - which additionally supports our decision to build a tool that covers all scenarios.

[15]The *planning* part was already described earlier in subsection III.b. The *managing* part comprises the following activities:

- Monitor and measure energy consumption and other costs constantly or in regular intervals - use detailed visualizations of power usage of servers and other hardware;
- If and when irregular patterns of energy usage occur, diagnose what is going wrong and where;
- Make necessary adaptations to counter the problem - allocate or cap power, improve efficiency of electrical and cooling systems in data center, improve service level automation, (re)evaluate: the electrical and mechanical support structures, architectural elements, data center's best practices, client's requirements for energy usage limit, etc.;

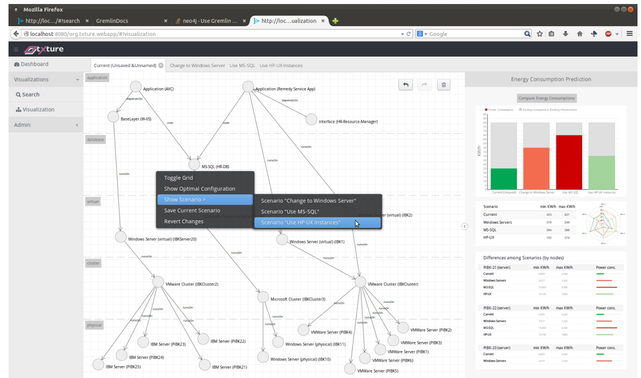

Figure 2. a) Switching between different configurations of the infrastructure, by selecting configuration from a right-click pop-up menu (central panel). The same can be achieved by clicking on the tab name in top of the window. b) Basic visualization (right-side panel) for different configurations, including comparison of KPIs, list of configurations with statistics on energy usage (min, max) and differences (in nodes/items) among configurations.

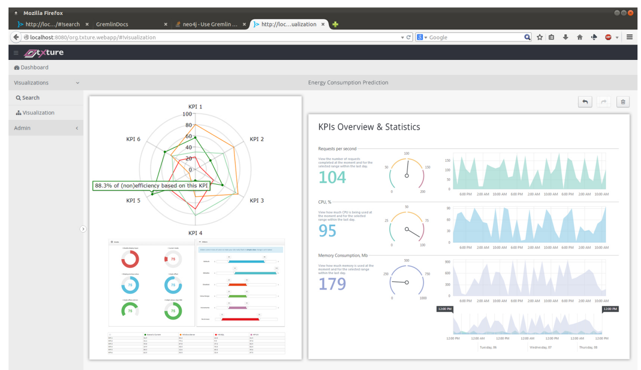

Figure 3. Advanced interactive visualizations, including radar graph and various statistical charts, of KPIs for different configurations.

- Ability to build and visualize *different configurations* (variations) of the same IT infrastructure and to show differences (in sense of architectural nodes) between them. This is done by starting from the base infrastructure configuration and building new ones with replaced items. User can view all configurations (figure 2.a).
- Ability to calculate, compare and visualize energy consumption and other performance indicators for different infrastructural configurations, based on available information[16], so as to allow the user to easily comprehend differences between these configurations (figures 2.b and 3). User chooses KPIs of interest from a set of predefined KPIs and can set minimum, maximum and weight values for each KPI, to better reflect his requirements. User can then adjust these values, and eventually decide which infrastructural configuration suits his needs best.

- Forecast power usage in advance based on consumption history and produce trend data for any single physical system within the infrastructure (automate energy management and optimization).

[16]Such information include (lowest, average and peak) energy consumption of a hardware device, calculation power, etc. Nowadays all hardware equipment is delivered with certificates including all such information.

B. Representation of KPIs

One of the visualizations used in the txture energy optimization module for representing KPIs is the so-called *Radar (spider) graph*, as proposed in the Green Grid Productivity Indicator Tool [14]. It allows for easy and clear visual interpretation for the stakeholders to assess how well the available resources are utilized, are the business objectives achieved and to locate the potential areas for improvement[17].

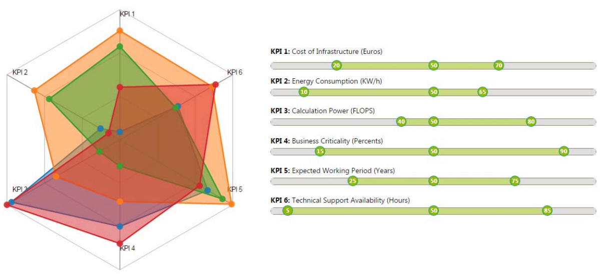

Figure 4. Radar graph with interactive selection (range) slider. Names of business values and corresponding KPIs (in this example each business value has exactly one attached KPI) are shown above selector bars. In this example, the left and right handle (min, max values) on the selectors can take values from 0-100. In real-world scenario, these numbers would take values according to possible business values. For example, cost of infrastructure would be represented in real currency, energy consumption in KW/h, etc. Weights always range between 0 and 100.

Multiple performance indicators can be visualized (for *several* different configurations at the same time[18]), against desired targets. As nicely explained in [5], target KPIs can be interpreted as business value parameters. Each business value parameter of interest to the stakeholder should have at least one, or more where applicable, metrics in the radar graph. Each axis[19] of the example radar graph shown in Figure 4 represents a performance indicator defined to have a value range of 0%-100%[20], where these two boundaries represent theoretical minimum and maximum values[21] for the KPI, ie. the least and most desirable values of the KPI. Performance indicators represented on a graph may or may not be related, as is with any metric that has several components. Selecting indicators with interdependencies and relationships is strongly advised, however - it can be very effective for holistic impact analysis. This also implies that a particular system change may result in some actual points becoming worse while others improve[22]. The radar graph offers several benefits, by providing a combined visualization of all performance metrics. It allows IT planners and operators to see the effects on all metrics simultaneously. It can be used to show existing state, change progress, and the impact of "what if" scenarios, as in different configurations of the same base IT infrastructure described in previous subsection.[23]

C. Outlook

The extensive research showed us both the significant potential of the concept as well as some space for improvement. In the subsequent iteration, we will extend our approach by implementing full automation of the optimization process.[24]

V. Validation of the Concepts and Project

In order to evaluate the proposed green EA model, framework and KPIs, a case study was performed. QE Research Group's[25] web servers, hosting a fictional project, have been compared, in order to find out which one is more energy efficient. One server uses 5 year old technology, while the other one uses new, modern technology. The previously described metrics are applied on both of the servers and results are then compared. The goal was to show that there

[17]By looking at a radar graph, the stakeholder can see which metric is the worst - this may be the metric with the greatest area for improvement.

[18]This is a clear advantage of our usage of radar graph in contrast to the way of usage proposed by Green Grid and found elsewhere. It allows for direct (visual) comparison of *several* different infrastructural configurations.

[19]Radar graph can have many axes, depending on the selection of metrics chosen by the stakeholder, up to a level allowing for clear viewing.

[20]In some cases, clear maximum values don't exist (e.g. for PUE). Then the axis end-points are established based on target values or other estimates.

[21]Several potential values of KPIs are of interest [5], [14]:

- Maximum and Minimum - 100% and 0%, as explained - best and worst values that can be (*theoretically*) achieved;
- Target value - A maximum (*in practice*) value considered a desirable and achievable goal (less than 100%);
- Actual value - A value, which might or might not be satisfying, calculated from real system data. This can be either the *peak value* recorded while monitoring, or the *average value* for the whole monitoring period.

[22]For example, improving server utilization may actually decrease UDC, which may lower DCiE if the data center doesn't scale with load. [14]

[23]First, a user can select KPIs of interest - those that will be represented in the visualization (radar graph). This is important because different enterprises have different interests and concerns; they care more about some and less about other KPIs (i.e. related business values). It is important to note that a single business value can have *several* related KPIs (for example, energy usage can be directly evaluated by PUE/DCiE, TUE, ITUE... and indirectly by many more KPIs), and thus users should carefully choose the KPIs of interest, to avoid redundancy. Then, users build different configurations of the starting (base) infrastructure. The tool will automatically calculate required KPIs for all these configurations and represent them on the radar graph. Each configuration is represented by a polygonal shape (outlined or solid, according to personal preference) of different colour. For example, Figure 4 shows a radar graph visualizing 4 different configurations previously built by the user, based on 6 chosen KPIs. This further means that better configurations (in terms of utilization) are those that have higher values of KPIs. Users can only rely on comparing configurations on one axis (by one KPI) at a time; however, even this can, in most cases, give hints towards the best configuration. Users can then interact with the radar graph by using the selection (range) sliders on side, as shown in Figure 5. Each selection slider represents one specific business value and is thus connected to one or more specific KPIs represented on the graph itself. Each slider has three handles which can be set to different values by user (click-and-drag). The left and right handles represent, respectively, minimum and maximum value for the corresponding business value. Meaning, the user can set the desired minimum and maximum values for the business value parameter, and the graph will be automatically updated to represent the changes in the value(s) of KPI(s) related to this business value parameter. The middle handle represents the weight of the business value. By moving it and setting it to a different value, the user can put more or less importance to the business value and corresponding performance indicator(s). This is important because enterprises put different "weights" to different business values. The same applies to involved stakeholders within the same enterprise - different stakeholders have different interests, priorities and benefits; thus, the ability to choose the KPIs (business values) and their max, min and weight values helps with reconciling their, sometimes different, even opposing, wishes. Allowing the setting of max, min and weights for KPIs enables immediate apprehension of the effect of potential changes which are planned for the data center and expressed through KPIs. E.g. a change at a data center can improve one, but deteriorate another performance metric at the same time. More examples of the radar graph's use can be found in [5], [6], [9], [13].

[24]This automation implies the ability of the energy optimization module to automatically calculate and propose the most optimized infrastructural configuration, for either complete or partial optimization of the infrastructure, on any chosen set of KPIs (with set minimum, maximum and weight business values), based on available information. For partial optimization, user can choose which architectural nodes and items can be optimized and which are not to be affected. This automation also implies that users will not have to create and compare different configurations of the base infrastructure, by using KPIs and interactive radar graph, thus saving valuable time and effort. However, users can still use radar graph (and its range selectors) to set min, max and weight for the business values.

[25]http://qe-informatik.uibk.ac.at - Quality Engineering (QE) is a research group within the Institute of Computer Science at University of Innsbruck.

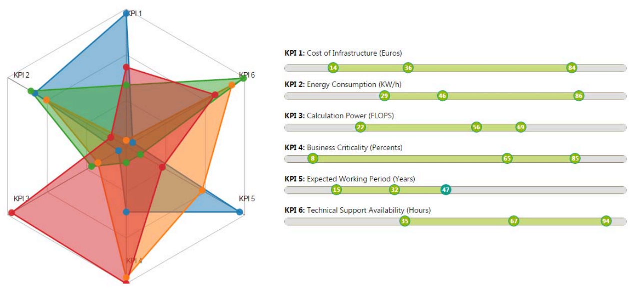

Figure 5. Radar graph with changed min, max and weight values for each business value.

are significant differences in energy efficiency, which can be identified by applying these metrics. The details on the case study can be found in the aforementioned Deliverable for EN-ACT project.

Furthermore, all the existing "green IT" KPIs were collected through an extensive literature review and then intensively compared and evaluated, with the aim of providing IT experts with enough metrics (within the energy optimization software) to monitor and optimize any aspect of the processes within the IT infrastructures. Evaluation was performed in several case studies[26], by testing the KPIs on IT infrastructure configurations of different sizes and complexities, to see which ones (KPIs) will give the most useful knowledge about the system, in terms of energy and other costs efficiency.[27]

In order to develop, test and evaluate the concepts and requirements for the energy optimization module, we have repeatedly consulted IT industry experts[28]. As a proof of concept, tests and case studies for evaluation of ideas, concepts, requirements and interface prototype have been conducted. After each new cycle of consulting IT infrastructure experts and incorporation of their suggestions into our research directions and concepts for the tool, we have conducted tests[29] of the updated concepts. To further test our approach, concepts and proposed solution, each of these tests was conducted based on different set of chosen KPIs.[30]

Eventually, we found out that our approach, with established concepts, worked fine for different scenarios.

[26]Some of which included assistance from IT industry experts, most notably from *QE LaB Business Services*: http://www.qe-lab.com

[27]Some of the key questions considered in these case studies were: How important different business values and KPIs were in their data centers, and which ones? Did their native enterprises implement any green initiatives, like *Green Grid* and *DPPE* (Datacenter Performance per Energy framework), and what were the results of those implementations?

[28]The meetings with industry IT experts proved to be very valuable, pointing out to the need for a semi- or fully-automated processing functionality of the energy management software. This need comes from the fact that often IT infrastructures consist of several thousands nodes - it is very cumbersome and time consuming to manually create and compare different configurations for such large infrastructures. Instead, the module should be able to automatically propose the best solution, based on previously selected KPIs (prioritized by importance), and taking into consideration available hardware solutions (equipment).

[29]These tests consisted of creating several IT infrastructure scenarios using txture tool, starting from a single IT infrastructure configuration - the "comparison specimen". We took two directions: scenarios of *similar* and *different* size, configuration and complexity. In addition, planning of a completely new IT infrastructure was also conducted. All three scenarios were used for comparison against the "comparison specimen". More details on the results of these tests will be published in a separate publication.

[30]This was needed to better emulate the real world industrial environment and different needs and considerations of different enterprises and stakeholders, as described earlier.

VI. Conclusion and Future Work

This paper sketched challenges and our approach - specialized software - to optimizing energy usage and other costs within large IT infrastructures, based on use of green KPIs and advanced visualization solutions. Our approach supports not only the *energy-aware* optimization of existing, but also *energy-aware* planning of new IT infrastructures. In the continuation of the project, we have started to prototypically implement our approach within the context of EA tool *txture*. More in-depth testing will be performed once the first working module prototype is available. The implementation and testing phases will be followed by additional evaluation activities and further improvements of the framework. Furthermore, we are continuing research of existing KPIs, in order to include as many as possible into our envisioned energy optimization module. This is particularly important in sense that no such thing as a "perfect" KPI exists - one that gives all the answers. To solve this problem, other KPIs need to be defined and available to IT planners and operators. We are intensively consulting IT experts on these recently identified KPIs as well, to make sure that their needs are fully met.

Acknowledgment: This research was part of the *EN-ACT (ENergy-Aware CompuTing)* project (www.en-act.eu) funded by *"Interreg IVA Italy-Austria 2007-2013"* program (www.interreg.net).

References

[1] M. Owen and R. Shields, "Using enterprise architecture to develop a blueprint for improving your energy efficiency and reducing your carbon footprint," 2008.

[2] I. Hanschke, *Strategic IT Management: A Toolkit for Enterprise Architecture Management*. Springer Science & Business Media, 2009.

[3] C. Hammond, "Introduction to green it," Fujitsu, 2012.

[4] N. Rasmussen, "Implementing energy efficient data centers," 2006.

[5] M. Dunbar and J. Espy, "Assessing and improving data center storage-related energy efficiency," 2008.

[6] L. Wang and S. U. Khan, "Review of performance metrics for green data centers: a taxonomy study," *The Journal of Supercomputing*, vol. 63, pp. 639–656, 2013.

[7] L. Newcombe, "Data centre energy efficiency metrics," Data Centre Specialist Group.

[8] C. Belady, "Green grid data center power efficiency metrics: Pue and dcie," The Green Grid, 2008.

[9] "Harmonizing global metrics for data center energy efficiency," The Green Grid, 2012.

[10] A. Kipp, T. Jiang, M. Fugini, and I. Salomie, "Layered green performance indicators," *Future Generation Computer Systems*, vol. 28, no. 2, pp. 478–489, 2012.

[11] S. S. Mahmoud, I. Ahmad *et al.*, "Green performance indicators for energy aware it systems: Survey and assessment," *Journal of Green Engineering*, vol. 3, no. 1, pp. 33–69, 2012.

[12] C. Cappiello, M. Fugini, A. M. Ferreira, P. Plebani, and M. Vitali, "Business process co-design for energy-aware adaptation," in *Intelligent Computer Communication and Processing (ICCP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 463–470.

[13] D. Cole, "Data center energy efficiency–looking beyond pue," *No Limits Software, White Paper*, vol. 4, 2011.

[14] C. Belady, "The green grid productivity indicator," The Green Grid, 2008.