**Association for Information Systems**
**AIS Electronic Library (AISeL)**

All Sprouts Content Sprouts

4-16-2011

# Application Driven IT Service Management for Energy Efficiency

Cinzia Cappiello
*Politecnico di Milano*, cappiello@elet.polimi.it

Alexandre Ferreira
*Politecnico di Milano*, ferreira@elet.polimi.it

Mariagrazia Fugini
*Politecnico di Milano*, fugini@elet.polimi.it

Barbara Pernici
*Politecnico di Milano*, pernici@elet.polimi.it

Follow this and additional works at: http://aisel.aisnet.org/sprouts_all

Recommended Citation

Cappiello, Cinzia; Ferreira, Alexandre; Fugini, Mariagrazia; and Pernici, Barbara, " Application Driven IT Service Management for Energy Efficiency" (2011). *All Sprouts Content*. 424.
http://aisel.aisnet.org/sprouts_all/424

This material is brought to you by the Sprouts at AIS Electronic Library (AISeL). It has been accepted for inclusion in All Sprouts Content by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

芽|Sprouts Working Papers on Information Systems ISSN 1535-6078

# Application Driven IT Service Management for Energy Efficiency

Cinzia Cappiello
Politecnico di Milano, Italy

Alexandre Ferreira
Politecnico di Milano, Italy

Mariagrazia Fugini
Politecnico di Milano, Italy

Barbara Pernici
Politecnico di Milano, Italy

Pierluigi Plebani
Politecnico di Milano, Italy

**Abstract**
Considering the ever increasing of information technology usage in our everyday life and the huge concentration of computational resources at remote service centers, energy costs become one of the biggest challenging issues for IT managers. Mechanisms to improve energy efficiency in service centers are divided at different levels which range from single components to the whole facility, considering both equipment and application issues. In this paper we focus on analyzing energy efficiency issues at the application level, focusing on e-business processes. Our approach proposes a new method to evaluate and to apply green adaptations strategies based on the service application characteristics with respect to the business process taking into account non-functional requirements.

**Keywords:** Green information systems, adaptive services, Green Performance Indicators

**Permanent URL:** http://sprouts.aisnet.org/10-131

**Copyright:** Creative Commons Attribution-Noncommercial-No Derivative Works License

**Reference:** Cappiello, C., Ferreira, A., Fugini, M.G., Pernici, B., Plebani, P. (2010). "Application Driven IT Service Management for Energy Efficiency," Proceedings > Proceedings of SIGGreen Workshop . *Sprouts: Working Papers on Information Systems*, 10(131). http://sprouts.aisnet.org/10-131

芽|Sprouts

## INTRODUCTION

The impact of IT on the energy consumption is continuously increasing and IT is becoming a major energy consumer. Such trend is due not only to the pervasive diffusion of information systems, but also to the availability of IT services of common use, such as for instance search engines, file sharing, and social networking services. Cloud computing paradigm is starting to provide the basis for new ways of providing services over rented resources from infrastructure providers (Zhang et al., 2010). While offering the possibility of providing services more easily and without the need of heavy infrastructure investments by service providers, these trends are increasing the need of managing efficiently large computer centers: energy consumption becomes an important item of cost and also a limitation if the power supply is not adequate for serving requests in peak times. With respect to that, a lot of work has been done at the infrastructure which includes IT hardware equipment and supply systems such as cooling and power system. On the other hand, our proposed work aims to tackle energy costs at the business process level such that we may need less physical resources, or at least use them in a more efficient way, to provide acceptable performance requirements.

## ADAPTIVE AND PROACTIVE SERVICE CENTER MANAGEMENT

In this paper we focus in particular on the proposal of managing resources at all levels in a service center based on two main approaches: a) a flexible management of resources, based on an intelligent information system for providing business process context-aware adaptive behavior in the data center; b) a deep knowledge of the services provided in the service center, to be able to anticipate their needs of resources and manage them in an energy-efficient way (e.g., Ardagna et al., 2008, Ferreira et al., 2009).

Regarding the former approach, we argue that the design of energy efficient information systems can leverage on some characteristics of the service-oriented approach and its flexibility and adaptivity, which can be applied at different levels and phases. In particular, while process analysis is usually focusing on performance, trade-offs between performance and energy consumption are also considered. Energy-related aspects include the amount and characteristics of energy resources which are available at a given time; peak situations may constrain the functioning of data centers causing an impact on business processes; the costs of different energy sources, especially at peak times; and the availability of computing resources for each assigned business task. Possible solutions may also involve rescheduling mechanisms when resources are not available or energy consumption is constrained. A self-* approach, as proposed in the service engineering area (Kephart et al., 2003), can be pursued, in particular to adapt systems to variable operating conditions and to different phases of the life-cycle of applications and services. Frameworks for adapting service-based business processes such as the PAWS framework (Ardagna et al., 2007) can be adopted to realize context-aware service systems.

However, in order to be able to provide such an adaptive behavior, we first need to understand all information regarding the service center and its running applications in details. For that reason, the monitoring, analysis, and control of the use of the service center components have to be supported by the development of a green controller which provides filtering and decision support functionalities in order to achieve pre-defined green strategies goals. Such strategies take into consideration energy efficiency at the service level in terms of

芽|Sprouts



(cc) BY-NC-ND Sprouts - http://sprouts.aisnet.org/10-131

Green Performance Indicators (GPI) which link low level measurements with high level business process features (Barroso, 2007).

For that reason, we have defined two groups of GPI: tactical and operational. GPIs belonging to the tactical group denote how the business process will consume less energy if its development phases are enhanced through the use of process engineering metrics. Such metrics characterize the level of maturity and capacities of the service development platform (including code language and development frameworks), the developer's use of innovative processes, and the rate of reused stable and well-known functions and methods. They are likely to produce energy-aware systems through improvement of the system quality in terms of service delivery versus customers' expectations and in terms of less complexity of the overall service interfaces.

In the operational group, we have GPIs that based on the information provided by the monitoring of IT resources, such as processor, memory, storage, etc., deal with energy needs and resource usage patterns from a business process perspective. To do so, GPIs abstract monitored values through some calculations in order to keep it independent from physical configurations and to be able to estimate business process needs in terms of power consumption (i.e. watts) and usage (i.e. percentage). The information provided by such GPIs is then used within a reasoning engine which is able to foresee possible situations in which the system can be set to a lower power mode by combining business process adaptation strategies and infrastructure allocation mechanisms such as dynamic workload scheduling. The application minimum quality of service requirements are used as thresholds to guarantee the application functionality also when lower power modes are adopted during execution.

The second direction mentioned in the beginning of this section is toward a proactive approach to the service center management, with the goal of continuous improvement of energy efficiency.

## RESEARCH QUESTIONS

Starting from the discussion introduced in (Watson et al., 2010), we propose a first analysis of a service center characteristics in Table 1, in which we discuss some of the proposed research questions considering service centers and different aspects to be considered when managing information for a more efficient use of computing resources.

In the table, we generalize the concept of sensor network to a more general concept of monitoring, the flow network in this case are service networks, sensitized objects are requested applications run in the service center. The information system in this case is an IS for assessment and control to drive the adaptive behavior of the service center.

In the first line we introduce issues about the level of granularity of the sensor network. In fact, such value depends on several other parameters such as facility dimensions and the type of the equipment or software that will be monitored. The biggest problem here is to find a balance between the overhead produced by the monitoring system and its benefits toward energy efficiency. It has a direct impact on RQ2 and RQ3 in which we discuss the needed granularity level of energy costs reduction. At RQ6 we point out the narrow correlation between supply and demand sides and at RQ9 we present examples that energy efficient strategies may bring both to customers and providers.

芽|Sprouts



(cc) BY-NC-ND Sprouts - http://sprouts.aisnet.org/10-131

## APPLICATION SCENARIO

In order to verify our proposed approach and demonstrate the advantages of the ideas proposed at Table 1, in this section we consider a generic e-commerce scenario in which some GPIs are considered to enact green adaptation strategies. The goal is not to present any specific e-commerce system, but to use some of the common activities shared by most of the e-commerce systems and their main characteristics. Let us assume that the system is composed by several individual tasks that are responsible to execute tasks such as customer registration, product selection, payment modules, and product shipment. In addition, each task is annotated with its minimum performance requirements that include physical resource considerations.

In order to perform such tasks, one or more application services are available with similar functionality but different from the non-functional point of view. Therefore, it enables us to select the best service maximizing performance, but, at the same time, minimizing energy consumption parameters. As explained in the previous section, applications are annotated with respect to relevant GPIs and performance values. Figure 1 depicts how our proposed approach works within this scenario. The GPIs selection phase selects among all existing GPIs the ones that are suitable for this specific application both at design-time and run-time. At design-time we have selected as a GPI *Application performance* that measures the number of transactions executed within one watt, and *IOPS/watt* which measures the number of I/O operations per second per watt. The GPIs selection for process design is mainly based on the nature of the application, which in this case performs a large number of database accesses. The GPIs selection phase also selects GPIs to be computed at run-time, i.e. during its execution, which are *CPU usage*, which measures the percentage of use of the CPU; *storage usage* which calculates divides the amount of used space by the total allocated space; and *energy consumption* which monitor the power consumption of each transaction. The last component, process feedback, is responsible to provide feedback information about GPIs used in order to improve future GPI selection phase.

| | **Monitoring** | **Service center** | **IS for assessment and control** | **Applications** |
|---|---|---|---|---|
| *RQ 1: What is the optimum level of information granularity of the sensor network to optimize a given flow network?* | Specific sensors placed at strategic points monitor physical infrastructure in terms of energy consumption, temperature, and humidity producing minimum possible overhead | Adaptive actions in the service center at different granularity levels based on different aggregated information | Evaluation of Green Performance Indicators (GPI) at different level of granularity taking into account each particular context | Consider functional and non-functional requirements (such as QoS) for applications and their components services, define optional services |
| *RQ 2: What information granularity enables effective enforcement of energy policy?* | Design the minimal monitoring infrastructure with less impact within the environment | Enable context-aware adaptivity actions | Dependencies among GPIs | Combine strategical and tactical goals |



(cc) BY-NC-ND Sprouts - http://sprouts.aisnet.org/10-131

| | | | | |
|---|---|---|---|---|
| *RQ 3: What information, and at what level of granularity, is required to optimize a given type of flow network?* | The monitoring system is composed by physical infrastructure (sensors) and logical components (application) | From physical devices, the granularity of the information starts from a single component (CPU) being aggregated until the track or room level | Context-aware energy management rules (at application, platform, infrastructure and facility levels) | Applications annotations to characterize energy consumption and utilization profiles |
| *RQ 4: What government policies and regulations will impel flow network managers to make them more energy efficient?* | The limitation here is the overhead produced by the different granularity levels of the monitoring system | E.g. apply EU code of conduct for data centers (EU code of conduct for data centers, 2010) | | Design application following software engineering best practices in order to create more energy-efficient codes |
| *RQ 6: How can an information system integrate supply and demand data to increase energy efficiency?* | In the monitoring system of a service center supply data is regarding to auxiliary systems such as cooling and power systems | Mechanisms to improve demand systems (e.g. IT hardware) will directly reduce supply side energy requirements | Manage resources at a global level, considering both performance and energy efficiency | Energy-aware service adaptation strategies (short-term) and service evolution (long-term) |
| *RQ 9: What information do consumers need about the usage of the objects they own or manage to increase their energy efficiency?* | | Reducing the infrastructure power consumption may lead to governmental benefits for both side (provider and consumer) | Energy profile for application | Mapping between needed resources for each application and available ones |

**Table 1.** Energy Informatics research challenges in service centers

Based on the values provided by the GPIs, different adaptation strategies can be applied both to the infrastructure layer and to the application layer. Regarding to the former one, such strategies are mainly oriented to changing physical devices mode, e.g. low power mode, turning off physical resources and virtual server migration. At the application level, adaptation strategies may include minimum performance requirements renegotiation, e.g. response time, by replacing one available application service by another one which is also available for the same task, and task skipping which can be applied over optional tasks only or reducing some of the functionalities. Considering our e-commerce scenario, the combination of *CPU usage, storage usage and energy consumption* GPIs at run-time may lead to virtual server migration in order to either provide less CPU power (it is mostly a database system) or stay closer to the required storage system aiming to reduce network overhead. By monitoring the usage parameters in the storage system we are able to play with parameters like data access frequency and data migration in order to change disk arrays mode or even turn them off. Furthermore, our tactical GPIs make sure that none of the selected service application is overtaking the number of transaction by useless operations.

芽|Sprouts



(cc) BY-NC-ND Sprouts - http://sprouts.aisnet.org/10-131

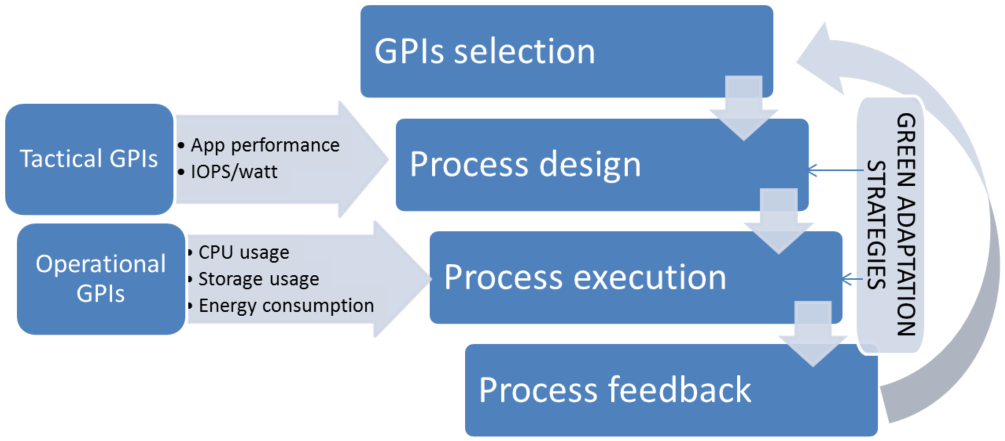


**Figure 1.** GPIs application over the methodology

## CONCLUDING REMARKS

In this paper we have presented our initial ideas to reduce energy consumption within service centers which focus at the application level at both design-time and run-time directions. Also, through a realist scenario example we provide details about its functionality and desired goals which embrace most of the issues with proposed solutions presented at Table 1 to tackle energy cost at the application level. The approach combines the high level view of business process and the low level of the infrastructure physical components. The approach presented here is an ongoing research work and first experimental results are just coming up.

## ACKNOWLEDGMENT

Part of this work has been funded by the European Commission within the Green Active Management of Energy in IT Service centers (GAMES) FP7 project.

## REFERENCES

Ardagna, D., Cappiello, C., Lovera, M., Pernici, B., Tanelli, M.: *Active energy-aware management of business-process based applications*. In: Proceedings of the ServiceWave. (2008)

Ardagna, D., Comuzzi, M., Mussi, E., Pernici, B., Plebani, P., *PAWS: A Framework for Executing Adaptive Web-Service Processes*, IEEE Software 24 (6) (2007) 39-46.

Barroso, L. A., Holzle, U. *The Case for Energy-Proportional Computing*. IEEE Computer Society Press, Computer 40 12 (2007), 33-37.

EU code of conduct for data centers, 2010. Available at: http://re.jrc.ec.europa.eu/energyefficiency/html/standby_initiative_data_centers.htm

芽|Sprouts



(cc) BY-NC-ND Sprouts - http://sprouts.aisnet.org/10-131

Ferreira, A.M., Kritikos, K., Pernici, B.: *Energy-aware design of service-based applications*. In: Proceedings of ICSOC/ServiceWave. (2009)

Kephart, J. O., Chess, D. M. *The Vision of Autonomic Computing*, IEEE Computer 36 (1) (2003) 41-50.

Watson, R. T., Boudreau, M.-C., & Chen, A. J. *Information Systems and Environmentally Sustainable Development: Energy Informatics and New Directions for the IS Community*. MIS Quarterly, 34(1), (2010), 23-38

Zhang Q., Cheng, L., Boutaba, R. *Cloud computing: state-of-the-art and research challenges*, Journal of Internet Services and Applications (2010) 1: 7–18

芽|Sprouts



(cc) BY-NC-ND Sprouts - http://sprouts.aisnet.org/10-131

芽|Sprouts
Working Papers on Information Systems | ISSN 1535-6078

*Editors*:
Michel Avital, University of Amsterdam
Kevin Crowston, Syracuse University

*Advisory Board*:
Kalle Lyytinen, Case Western Reserve University
Roger Clarke, Australian National University
Sue Conger, University of Dallas
Marco De Marco, Universita’ Cattolica di Milano
Guy Fitzgerald, Brunel University
Rudy Hirschheim, Louisiana State University
Blake Ives, University of Houston
Sirkka Jarvenpaa, University of Texas at Austin
John King, University of Michigan
Rik Maes, University of Amsterdam
Dan Robey, Georgia State University
Frantz Rowe, University of Nantes
Detmar Straub, Georgia State University
Richard T. Watson, University of Georgia
Ron Weber, Monash University
Kwok Kee Wei, City University of Hong Kong

*Sponsors*:
Association for Information Systems (AIS)
AIM
itAIS
Addis Ababa University, Ethiopia
American University, USA
Case Western Reserve University, USA
City University of Hong Kong, China
Copenhagen Business School, Denmark
Hanken School of Economics, Finland
Helsinki School of Economics, Finland
Indiana University, USA
Katholieke Universiteit Leuven, Belgium
Lancaster University, UK
Leeds Metropolitan University, UK
National University of Ireland Galway, Ireland
New York University, USA
Pennsylvania State University, USA
Pepperdine University, USA
Syracuse University, USA
University of Amsterdam, Netherlands
University of Dallas, USA
University of Georgia, USA
University of Groningen, Netherlands
University of Limerick, Ireland
University of Oslo, Norway
University of San Francisco, USA
University of Washington, USA
Victoria University of Wellington, New Zealand
Viktoria Institute, Sweden

*Editorial Board*:
Margunn Aanestad, University of Oslo
Steven Alter, University of San Francisco
Egon Berghout, University of Groningen
Bo-Christer Bjork, Hanken School of Economics
Tony Bryant, Leeds Metropolitan University
Erran Carmel, American University
Kieran Conboy, National U. of Ireland Galway
Jan Damsgaard, Copenhagen Business School
Robert Davison, City University of Hong Kong
Guido Dedene, Katholieke Universiteit Leuven
Alan Dennis, Indiana University
Brian Fitzgerald, University of Limerick
Ole Hanseth, University of Oslo
Ola Henfridsson, Viktoria Institute
Sid Huff, Victoria University of Wellington
Ard Huizing, University of Amsterdam
Lucas Introna, Lancaster University
Panos Ipeirotis, New York University
Robert Mason, University of Washington
John Mooney, Pepperdine University
Steve Sawyer, Pennsylvania State University
Virpi Tuunainen, Helsinki School of Economics
Francesco Virili, Universita' degli Studi di Cassino

*Managing Editor*:
Bas Smit, University of Amsterdam

*Office*:
Sprouts
University of Amsterdam
Roetersstraat 11, Room E 2.74
1018 WB Amsterdam, Netherlands
Email: admin@sprouts.aisnet.org

芽|Sprouts