

Green IT and Data Center Efficiency

Antonio Paolacci

June 20, 2011

Abstract

The design of an energy efficiency data center requires attention, skills, experience and investments but if design activity is well done we have considerable advantages in future costs and management. Each component should not to be selected independently, but according to others and to the entire structure. Important factors to design efficiently are: to take account of internal flows of electricity and heat which are the main cost factors, to reduce greenhouse gas emissions, to avoid unnecessary oversizing of the IT structure, to maximize IT productivity and to choose high energy efficient IT equipments. All this concepts will be detailed below in this report.

Contents

1	Introduction	1
2	Key concepts	3
3	Best practices	8
4	Alternative energy	12

1 Introduction

The term green computing or green IT appears for the first time in the article *Harnessing Green IT: Principles and Practices*. The author San Mурugesan defines the field of study of green IT as “the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems, such as monitors, printers, storage devices, networking and communications systems, efficiently and effectively with minimal or

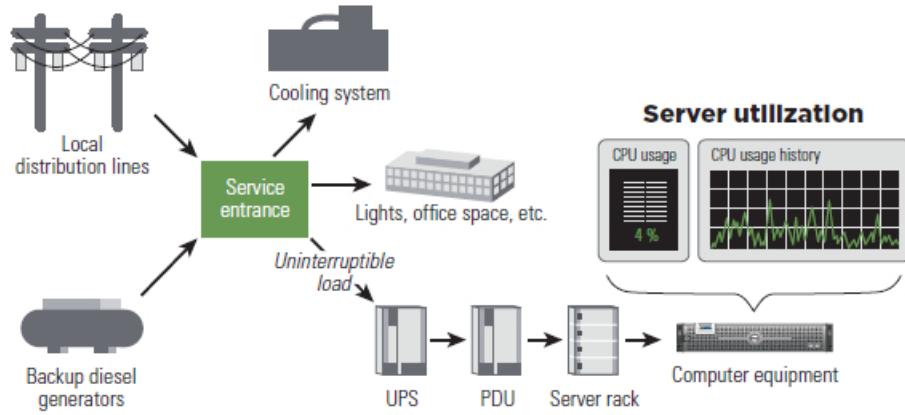


Figure 1: Power distribution to components of a data center

no impact on the environment.” The key points of this definition are the *environment* term, we can try to reduce the environmental impact and the release of pollutants(CO_2); and the *energy efficiency* term, that is to reduce the amount of energy(the Figure 1 shows energy distribution to data center components) required to provide products and services for IT industry. The author titles the paper *Harnessing...* expressing its intention to exploit benefits developping green data center. In fact energy efficiency is an environmental and IT economic imperative that couldn’t be ignored. Slowing down the growth rate of electric power use for server farms and related components is a necessary part of slowing down the rate of greenhouse gas emissions. CO_2 emissions from the use of fossil fuels (coal in large part) are leading contributor to global warming and climate change. Until cleaner plant technologies or some other alternatives become available, balancing electricity consumption with IT productivity of a data center is the most important thing that future policy must keep in mind. In the U.S. an EPA report and related opinion by Uptime Istitute¹ indicate that US servers alone consume 1.2 percent of the nation’s electricity in 2005, more than double the use over 2000 and equivalent to 38 million metric tons of CO_2 emissions. This is not a momentary blip, but is a long-term and growth problem with no sign of relevant abating. This level make IT hardware and infrastructure the leading contributor to carbon emission in US. A study by a physicist² at Harvard University,

¹ EPA, Enviroment Protection Agency. Uptime Institute’s published opinion titled “EPA Report Should Spur Industry-Wide Green Data Center Movement,” by Kenneth G. Brill, the Institute’s founder and executive director

²From The Sunday Times January 11, 2009 - Revealed: the environmental impact of Google searches

estimated that the average Google search released 7 grams of carbon dioxide (CO_2) as boiling a kettle for a cup of tea. Imagine millions of people type-in Google form without considering the environment. Now we can only try to research into key areas such as making the use of computers and servers as energy-efficient as possible, and designing algorithms and systems for green technologies, because is true that with energy saving we reduce the CO_2 emissions and reduce the global warming, but it's more true that we save lots of money. The shock billing of one of the best italian banking group is about 150 millions of euro/year equally shared between cooling, lighting, IT hardware; their pc's consume about 50 million/year, so every reduction point percent can earn 500 thousands euro. The adoption of virtualization techniques, also by SME's (Small and Medium Enterprises), allow to save money and increase ROI (Return Of Investement) index. Many people and firms think, incorrectly, that the green economy is much expensive than normal economic and less performant, and they choose it only for marketing reason: exhibit green computing certifications and logos on their business websites. Gartner US Research claims ³ that during the next five years, most US enterprise data centers will spend as much on energy (power, cooling) as they do on hardware infrastructure. Gartner estimates that more than 70% of global data center facilities must be modified during the next five years. The imperative is to 'change!' and many firms adopt a new key job figure: the eco-responsability expert, like David Douglas, eco-responsability di Sun Microsystems, that developped three new data centers in Europe, India, US and promises to reduce energy consumption of 60% than in the past. Siting of a data center is one of the factors that affect the energy consumption and environmental effects and in the near green future will be very important. In areas where climate favors cooling and lots of renewable electricity is available the environmental effects will be more moderate. Thus countries with favorable conditions, such as Finland, Sweden and Switzerland, are trying to attract cloud computing data centers.

2 Key concepts

In the following section are presented key concepts that should be followed like guidelines for the future of green IT. Some are technical aspects (benchmarking units or how to choose best IT components), some are simply exhortations.

³Analysts Examine the Latest Issues Related to Data Centers and Green IT During the Upcoming Gartner Symposium/ITxpo, October 7-12, in Orlando

- **Reset research and development priorities to include energy efficiency** The Research and Development people are now called to reset long-term developmental roadmaps to ensure taking attention to the problem of energy consumption of chip,server, network and storage devices. We need to slow-down the rate of increase in computational power according to the rate of improvements in electric consumption. We need to do so to restore the benefits of *Moore's law*⁴ on software and hardware. We can think to design intelligent algorithms able to move the most intensive computation on a server where the cooling is more efficient⁵. Then all IT vendors are called to make energy-efficiency a key feature of their customers information and education. Due to their influence if customers demand energy-efficiency improvements, vendors will provide it.
- **Benchmark Data Center Energy Consumption and Efficiency** What does *efficiency* really means in the context of a data center? Benchmarking data center's energy efficiency is a first step towards reducing power consumption and related energy costs. Benchmark techniques enable to understand the current level of efficiency in a data center but also verify the effectiveness of an additional implemented improvements by measuring the situation before and after. In a perfect information world each "choose" must be measured related to its power consumption: the use of a particular device rather than others, the use of a particular DBMS rather than others, the use of a particular protocol rather than others, the UPS⁶ system, the HVAC⁷ system,... There are many well-known metrics for data center energy costs. This report presents only the two important. The first, Corporate Average Datacenter Efficiency (CADE), was developed by the Uptime Institute. The institute determines the facility efficiency by calculating the amount of space used and multiplying that by the energy efficiency of the building. The higher the CADE number, the more efficient the data center.

$$CADE = Facility\ Efficiency(FE) * Energy\ Efficiency(EE)$$

$$FE = Facility\ Energy\ Efficiency * Facility\ Utilization$$

⁴The number of transistors that can be placed inexpensively on an integrated circuit doubles approximately every eighteen months.

⁵Energ-It project by Fondazione Politecnico and Dipartimento di Elettronica e Informazione del Politecnico di Milano

⁶UPS, Uninterruptible Power Supply

⁷HVAC, Heating Ventilation and Air Conditioning

$$EE = IT\ Energy\ Efficiency * IT\ Utilization$$

The other method is called Power Usage Effectiveness (PUE), developed by the Green Grid ⁸. In this method, the lower the PUE the better is. The components of the PUE calculation are “Total Facility Power” (TFP) and “IT Equipment Power” (IEP). TFP is measured at the utility meter for the data center space and includes all of the components required to support the IT load including: power components including UPS, cooling elements, other infrastructure components such as lighting. As might be expected, IEP, is the sum total of the power used by the facility’s computing components including servers, storage devices and networking devices. To calculate a facility’s PUE simply divide the Total Facility Power by its IT Equipment Power.

$$PUE = \frac{Total\ Facility\ Power}{IT\ Equipment\ Power}$$

PUE score of 2 indicates that for every watt of IT power, an additional watt is consumed to cool and distribute power to the IT equipment. A PUE of 1 would represent theoretically perfect efficiency, in which all energy was used directly by IT equipment. The lowest published Google PUE is 1.09 and maybe is the best in the world. There are many various metrics and there is no standard in this field so some people claim that benchmarking is useful only for example if you are comparing yourself with others with similar power and cooling equipment types, percentages of load utilization, energy costs.

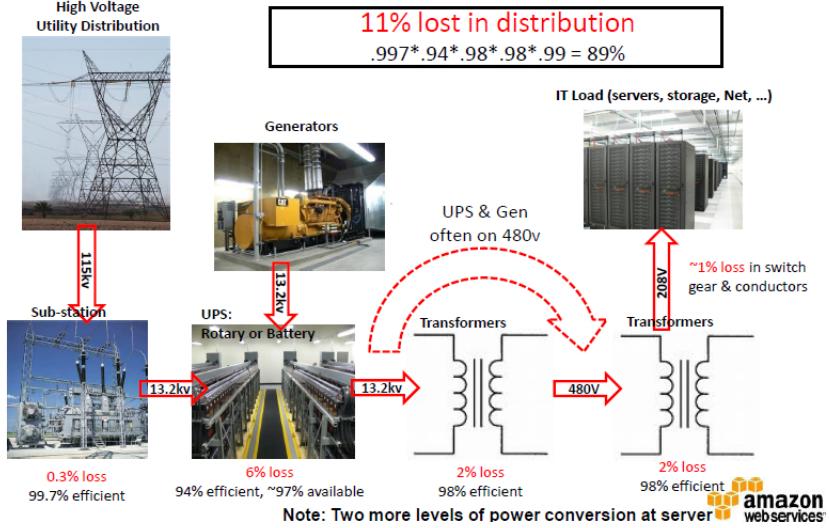
- **Adopt new IT governance policies** Regulation, law, social pressures, rising energy costs will bring data center industry to answer issues related to energy efficiency. For example in other field, the recent adoption of an unique standard for mobile-phone-charger or an unique regulation of stand-by consumption of digital components. And then IT vendors must include the true costs of site infrastructure management in the total-cost-ownership of IT hardware and business online applications.
- **Identify key features of a green data center** The optimization problem of a data center is difficult to correctly solve if we don’t understand on which elements act and what efficiency we want. The problem could be divided in four small subproblems like:

⁸The Green Grid, 2007, “The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE,” Technical Committee White Paper

1. *Maximize the fraction of IT equipment productively and fully utilized.* Consolidation and virtualization(see following) techniques could help in this phase for a good allocation of the resources. Could be necessary turning off comatose servers, servers which are running applications no longer needed or which run no application at all, and enabling server power-save features.
2. *Select hardware who has the best computational performance per Watt.* **Operations per Second/W**, it's a metric measured for many servers by SPEC (Standard Performance Evaluation Corporation), that collect server benchmark and makes them available on-line. We need to analyze each component evaluating, for example, economic indexes like Life Cycle Cost Analysis (LCCA) or Total Cost of Ownership (TCO). The US EPA has an **ENERGY STAR** program for large data centers, PC, server, monitor and printer, while in the case of the power supplies is available to consider **80PLUS** program. The Uptime Institute suggests to use efficient power supplies that convert AC power at the plug to DC power, used by many internal components.
3. *Select IT equipment who delivers input power to IT hardware most efficiently.* Select high energy efficiency UPS or PDU⁹. More than 20% of the energy is lost in the transformation process AC/DC in the PDU units and DC-transformation is considered by many as one of the power-consumption-reduction methods for data center environments. There's considerable disagreement over the potential energy savings afforded by switching from AC to DC powered servers and storage in the data center. The IT industry has been using DC-powered equipment for years, so there is precedent for going this route, but significant savings must be demonstrated to propel DC power into the enterprise data center. The driving force behind the switch is the inefficiency of today's AC power supplies. The typical AC to DC power supply found in most servers is 70% to 80% efficient, which means that 20% to 30% of every watt delivered to a piece of equipment produces nothing but heat. So you pay for the wasted electricity to run the equipment and also you pay for the electricity to cool the heat produced by that wasted electricity. In a data center using DC power, a rectifier-equipment that converts your utility's AC grid power to DC power-typically operates at over 90% efficiency and doesn't need to be located in

⁹PDU, Power Distribution Unit is an electrical device used to control the distribution of power to individual loads.

the data center. Recently Google invested in this area and buy *Transphorm*, an US company experimenting this innovation in data center.



Design well-sized cooling system: variable speed fans, free-cooling system, chiller¹⁰ with elevate EER(Energy Efficiency Ratio).

4. *Maximize the amount of useful power delivered to the site infrastructure at utility meter.*¹¹ Very soon the site infrastructure of a data center hosts also other services, it's important to install ad-hoc utility meter and monitore them.

- **Energy and Power are different concepts, so takes attention using these terms for calculations!** *Power* is a spot measurement at a particular point in time, *energy* is consumption in a period of time. Power is measured in kilowatts (KW) while energy is measured in kilowatts hours (KWh). There is a power measurement for every instant but users tipically care about of the “peak”, the highest power draw at any point of time. An IT manager that hibernates idle servers at night will save *energy* but not decrease the peak *power* demand. The site infrastructure will be sized accoording to the peak demand. If a data center really consumes the same amount of power every hour of the year the the energy used in KWh is the power in KW times 8760

¹⁰Chiller is a machine that removes heat from a liquid via a vapor-compression or absorption refrigeration cycle. This liquid can then be circulated through a heat exchanger to cool air or equipment as required

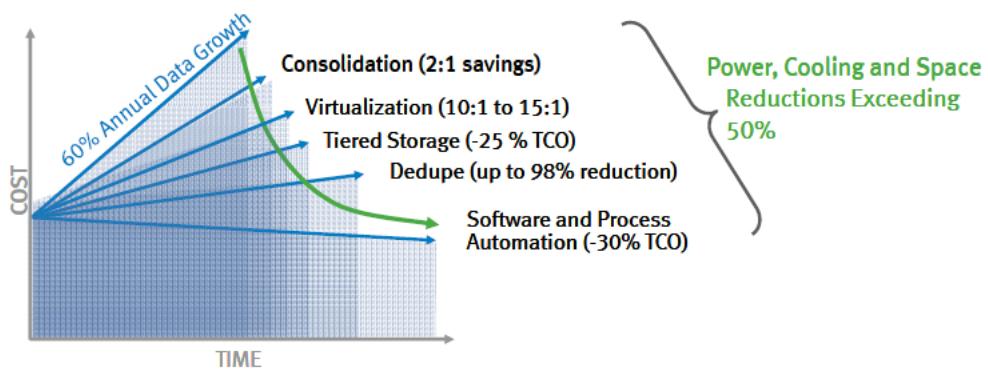
¹¹An “electricity” utility meter is a device that measures the amount of electric energy consumed by a residence, business, or an electrically powered device

hours. However in a green data center the annual energy use will not be the peak power times 8760 hours because:

- idle server and storage could be hibernated at off-peak time, at night or in the weekend;
- free cooling for some hours could be used in cold seasons of the year;
- future IT equipment could be designed to use less than the peak power when it is used for partial compute load.

3 Best practices

Many data centers, built 10-15 years ago, are underpowered and lack the cooling capacity to fully support today's high-performance, high-density servers produced by vendors. But despite their insufficient energy and cooling systems they are still usable with a good investment: evaluating actual infrastructure(space,power consumption,server room temperature,cables) and using tools such as virtualization to bridge the gap between IT load and other facility's components. In this section is presented what is done in practice in real word of data center. These tips are used to reengineering an old data center or to design a new by IT engineers (let see the following figure show expected benefits from these best practices).



Consolidation One of the most evident problems of a data-center is the increasing amount of servers, that increases without any form of control the costs of infrastructure. Server consolidation try to reduce the total number of servers. In fact in the past to activate a new service the IT managers need to install a new server machine. The potential benefits of a consolidation phase are:

- Reduced costs: less servers, less software licenses, less server location spaces, less energy power,...
- Improved control and manageability of server: “server room” more easy to manage.
- Improved service level: increased possibility to focus attention on higher value tasks and services.

Virtualization Server virtualization techniques allows to better share and use computational and storage power resources of a single server among plus services or applications. Start by eliminating the old style “one server, one app” model and run multiple virtual machine on each physical machine. Virtual machine can run different operating system and multiple applications on the same physical server thanks to a layer of software called “Hypervisor”. Virtualization reduces IT costs while increases energy efficiency, utilization and *flexibility*. Operational flexibility means to have the capacity for faster react to market changes with dynamic resources management or dynamic resources allocation. Consolidation and virtualization resources typically save 60% of IT costs.

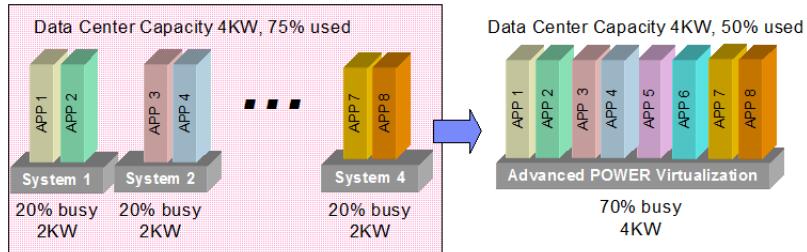


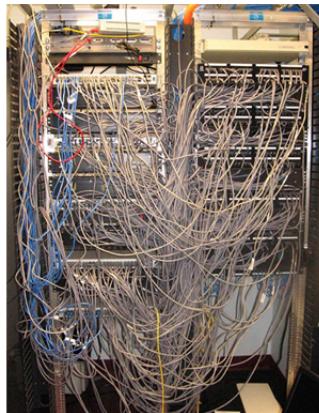
Figure 2: In the first case, without virtualization, total power is 8KW while in second case total power is 4KW

Tiered storage Tiered storage is a data storage environment consisting of two or more kinds of storage delineated by differences of this attributes: *price, performance, capacity and function*. Flash drives and disk compose two separate storage tiers. Old and new technology disk configured to use different functions such as RAID level or different data-access rate compose a two separate storage tiers. The classification of storage permits a better and efficient management and can reduce disk usage and power consumption if data are allocated on different ranking storage related to data access rate, data

transfer rate and data dimension.

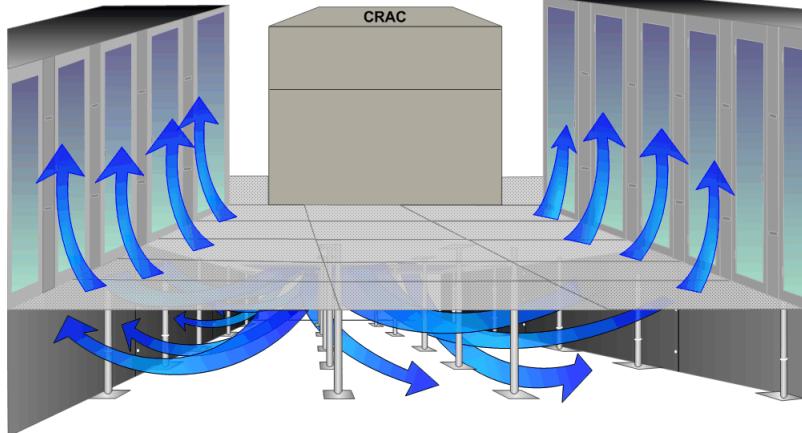
Dedupe The deduplication technology is able to find and eliminate redundant data. Some operation like backup or attach file to email sent to many recipients create a large number of duplicate data. Eliminating redundant data can significantly reduce storage number unit and improve bandwidth efficiency.

Structured Cabling Structured cabling divides the entire IT infrastructure into manageable blocks and then attempts to integrate these blocks to produce the high-performance and reliable network. Many IT service failure are due to an unstructured cabling and this kind of failure is very difficult to detect. Structured cabling contribute to realize a green IT infrastructure because it makes servers and switches to work with optimal performance. Cables blocking air inlets and exits will raise the temperature of server and related components and reduce the reliability. Moreover, unstructured cabling makes difficult to maintain servers and switches when there is a problem or if we want to move the rack in the server room.



Raised floor and cooling Raised floors and cooling system related to them, are another suggested solution to the problem of cabling the IT structure and cooling server-racks and others data center components. Note that the best solution maybe could be an hybrid between solid and raised floor. In a raised floor environments servers are typically cooled by pushing warm air away through high pipelines and blowing cold air under the floor. This process leverages the natural behavior of hot air to rise. It is significantly more difficult to cool a server environment by pushing cold air downward. In such case a raised floor can help. Then another benefit can be: routing infrastructure under a raised floor keeps hundreds or thousands of associated power cables and network cables out of sight, which makes them less susceptible to

damage or being unplugged accidentally.



Redundancy Redundancy is a key feature for a data center because the infrastructure, by its nature, need to work 24 hours a day, 7 days a week. The U.S. TIA (Telecommunications Industry Association) proposed standards TIA-942 in 1996 which have been adopted now at international level, called **TIER I, II, III e IV**. The survey is based on redundancy of energy distribution lines and equipments. To choose which is the TIER to adopt we take account of the real purpose and service level agreement(SLA) with the user to avoid unnecessary oversized of IT infrastructure and unnecessary energy consumption.

Tier I	Tier II	Tier III	Tier IV
only one path for power and cooling	only one path for power and cooling	multiple paths for power and cooling (only one activated)	multiple active power and cooling paths
no-redundant components	redundant components	redundant components	redundant components
no-raised floor	raised floor	raised floor	raised floor
single pathway for all telecommunication cabling	routers and switches have redundant power supplies and processors	be served by at least 2 access providers, redundant backbone pathways, multiple routers and switches for redundancy	backbone cabling should be in conduit or have interlocking armor, optional secondary distribution area, optional redundant horizontal cabling
downtime per year: 28,8 hours	downtime per year: 22 hours	downtime per year: 1,6 hours	downtime per year: 0,4 hours

4 Alternative energy

To cope with the change among green and old data center, some companies are looking to alternative energy sources: mainly to solar panels, natural gas turbines, wind energy, fuel cells, hydropower. For example in US, for the NCTD public transport company in San Diego, solar power was the obvious choice. Thirty solar panels were activated and has been virtualized the IT infrastructure. Solar energy obtained not cover entire energy needs of the data center. For this reason, the NCTD sells electricity and buy power for data center. The Syracuse University in New York has invested nearly 12 million of dollars to build a data center powered by natural gas turbines. In this case, contrary to what happens to NCTD, gas turbines cover all needs of the data center(1 MW) and at the same time providing 200KW of energy for other uses to the buildings of Syracuse University. Another solution for powering the data center is the use of **fuel cells**, which do not cause harmful emissions. Currently using this technology, large companies like Google, while medium and small enterprises can not use it because of the costs of building such facilities. Fuel cells provide energy from the reaction between a fuel supply and an oxidizing agent, that could be hydrogen(still not from H_2O) and oxygen(usually from air), producing H_2O . Another new solution was developed by Google with a new data center in Finland.[10] The plant will take the water of the Baltic Sea to cool server, which will be the first data center to be cooled in a natural way, through continuous exchanges of sea water will replace the chillers. However the final goals are always the same: to save considerable sums of money in the maintenance of a data center, to reduce the more expensive and problematic voice of a data center: the cooling and to mitigate enviromental impact.

Conclusion

Increasing data center energy efficiency is not only an environmentally friendly strategy, but also a key way to cut costs. By consolidating systems using virtualization, energy-smart servers and components, alternative energies and small useful practices, organizations can significantly reduce power and cooling requirements and create a green data center.

References

- [1] Kenneth G.Brill, founder and director Uptime Institute, “EPA Report Should Spur Industry-Wide Green Data Center Movement”, Opinion of

the Institute

- [2] Giampiero Calli Ballola, “Green Data center: percorsi, tecniche e idee”
- [3] “How To Benchmark Data Center Energy Costs”,
<http://www.environmentalleader.com/2009/01/13/how-to-benchmark-data-center-energy-costs/>
- [4] “Power Usage Effectiveness (PUE)” -
<https://www.digitalrealtytrust.com/pue-efficiency.aspx>
- [5] “Data Center Efficiency”, Slides of EMC Consulting Seminar, 2011
- [6] “Google Data Centers” <http://www.google.com/corporate/datacenter/index.html>
- [7] M. Bramucci,D. Di Santo,D. Forni, “Linee guida per la progettazione di datacenter ad alta efficienza”, ENEA Agenzia Nazionale per le Nuove tecnologie, l’Energia e lo Sviluppo Economico Sostenibile
- [8] P. Rad, M. Thoene, T. Webb DELL “Best Practices for Increasing Data Center Energy Efficiency”
- [9] Uptime Istitute, “Four Metrics Define Data Center “Greenness”
- [10] 24 Maggio 2011 Repubblica.it - Energia, l’efficienza conviene Google fa scuola alle aziende
http://www.repubblica.it/ambiente/2011/05/24/news/google_efficienza_energetica-16696225/