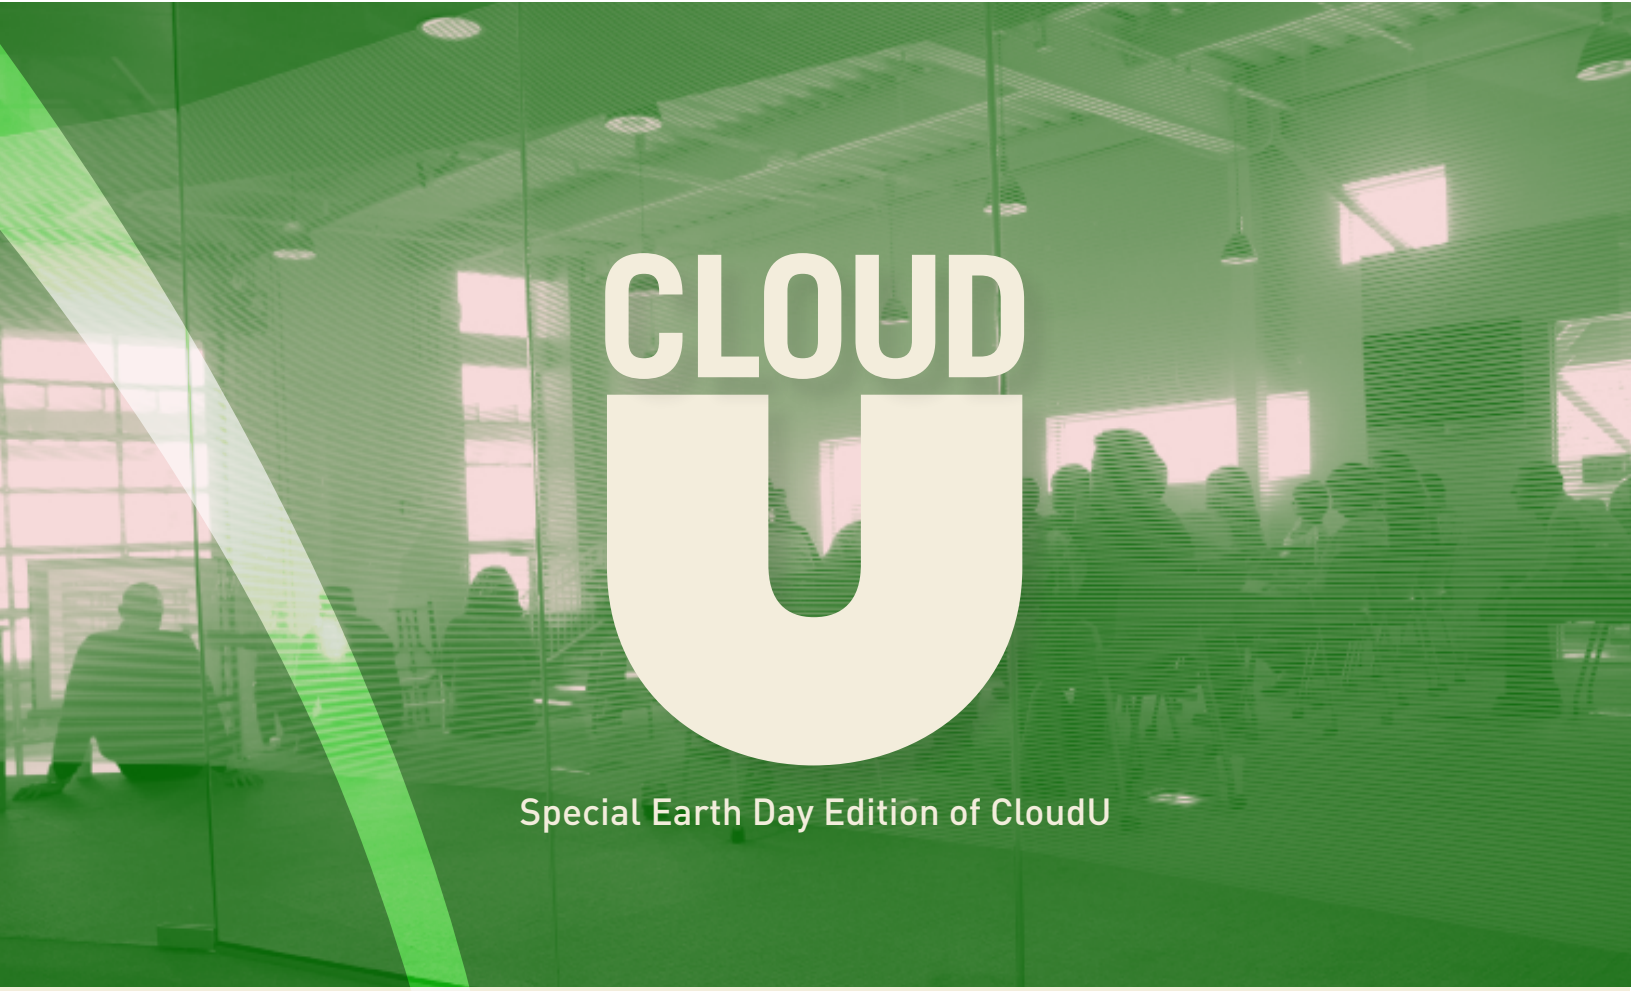

SUSTAINABILITY AND THE CLOUD

THE GLOBAL ENVIRONMENTAL BENEFITS OF IT HOSTING AND CLOUD TECHNOLOGY

Through this year-long series of whitepapers and webinars, independent analyst Ben Kepes will be building a cloud computing curriculum designed for technologists and non-technical users alike. The mission is to build widespread knowledge about the cloud revolution and encourage discussion about the cloud's benefits for businesses of all sizes. Read more CloudU whitepapers and register for upcoming webinars at www.rackspace.com/cloud/cloudu.

Diversity Limited

Sponsored By: rackspace® HOSTING

Diversity Limited

Table of Contents

Executive Summary

Earth Day is an annual event designed to increase awareness of our impact on the earth. Earth Day, first held in 1970 was founded on the premise that

> *...all people, regardless of race, gender, income, or geography, have a moral right to a healthy, sustainable environment.* [1]

Climate change is one of the most significant problems facing the world and impacts directly on weather patterns and indirectly on almost all aspects of our lives.

Earth Day strives to encourage people to think about their impact on the earth and how it can be reduced. To mark Earth Day 2011, we are publishing this special CloudU report to detail how Cloud Computing can aid in reducing Green House Gasses (GHG).

We contend however that there needs to be an economic imperative to change behaviour and, in the case of Cloud Computing, there are both moral reasons for a shift, but also sensible economic reasons to move from the traditional model of computing to Cloud Computing.

We detail how Cloud Computing, when compared to traditional computing, has multiple positive impacts both economically and environmentally. These benefits are driven by the tendency of Cloud Computing to;

- Reduce total infrastructure allocation
- Increase efficiency by leveraging multi tenancy
- Maximize server utilization rates
- Improve data center efficiency

We will detail how Cloud Computing drives these benefits and why we believe it is the right delivery method to minimize the impact of IT on the planet.

The Drive for Sustainability

Climate change is one of the biggest threats humankind will need to face in the next century. While debate rages about the finer details of climate change, projections suggest [2] that;

- Decreased Air Quality due to increased ground level ozone and particulate matter will impact human health
- Sea levels may rise as much as 59 cm (23 inches) during the 21st Century
- Melting ice may lead to changes in ocean circulation and speed up warming
- Changes in weather patterns lead to increased extreme weather events negatively impacting human populations and economies.
- Changes in habitats will lead to an increase in species extinction

In the face of such calamitous effects, we all need to look at ways of reducing our impact on the earth. The internet and computing are now an integral part of our daily lives. The ability to learn, communicate and shop online exchanges one set of resources (cars, fossil fuels) with another (computers, the Internet, servers) but creates another set of challanges to overcome. Cloud Computing is one of a range of innovations that can reduce our footprint, and in a meaningful way improve an organization's bottom line. The uptake of Cloud Computing is being driven by a number of factors including;

- The modern demand for computing to be increasingly mobile and less expensive
- Organizations' requirements to move from capital expenditure to operating expenditure
- The more volatile nature of commerce and business generally
- Cost reductions across the board

At the same time we are seeing more awareness around the impacts of our consumption generally, and IT specifically. This comes at a time when there is upwards pressure on the cost of electricity driven by;

- Resource constraints
- Ever increasing demand leading to higher unit costs
- Concerns over the environmental impacts of electricity generation
- The need to invest in electricity infrastructure to continue to supply the increasing demand
- Legistlation to improve air quality and regulate utilities

Concurrently with all of this there is a move from individuals and organizations to take more of an interest in the environmental sustainability of their business – for both social and profit reasons.

The intersection of all of these factors has impacts on the way computing is delivered, and we contend it will further drive the move towards Cloud Computing as opposed to a traditional approach with on-site IT infrastructure.

The Current Situation

In the face of increasing calls to reduce greenhouse gas (GHG) emissions in an effort to decrease climate change, it is important to reflect that the Information and Communications Technology (ICT) "industry, as of 2007 accounted for 2% of total global carbon emissions [3]. This figure is set to rise to 15 percent of total global carbon emissions (or 7.8 billion tons of CO2 equivalents per annum) by 2020 [4]. Reports suggest [5] that the environmental footprint from data centers is set to more than triple between 2002 and 2020. This result would see data centers become the fastest-growing contributor to the total carbon footprint of the technology sector.

As the ICT sector continues to grow towards a predicted 8.7% of global GDP by 2020 [4], innovations within the sector will both increase efficiency and innovation, while reducing the relative impacts of the IT industry itself on global carbon emissions. The bulk of business computing will shift out of private data centers to the 'cloud.' [6] Cloud Computing with its shared resource model, tends to be significantly better utilized, and more efficient, than traditional on-premise data centers. The following pages will demonstrate why.

Doing Well By Doing Good: Business Benefits Tied to Environmental Benefits

In difficult economic times, organizations are reluctant to change behaviors for merely altruistic reasons, they require a real business benefit, be it cost-reduction or revenue increase, in order to change the way they work. For this reason, Cloud Computing with the economic benefits it drives is uniquely placed as both a driver of sustainability and a driver of business performance.

In a recent report [7], Gartner stated that the move to Green IT and sustainability generally is about business value. Gartner stated that;

> *Pursuing an integrated and information-enabled sustainability strategy protects your organization from current and future regulatory requirements, solidifies reputational value and provides access to substantial savings through radical, technology-enabled efficiency gains and a variety of capital.*

We believe that Cloud Computing, and in particular some of the benefits that virtualization can drive, meets Gartner's requirements in terms of sustainability and business value and we will now take a quick look at what Cloud Computing is and how it drives these benefits.

Cloud Computing In Brief

For a detailed explanation of what Cloud Computing is, and how it creates a revolution within IT, refer to the CloudU report "Revolution Not Evolution - How Cloud Computing Differs from Traditional IT and Why it Matters" [8] For the purposes of this report, we will simply use the National Institute for Standards and Technology (NIST) definition of Cloud Computing to frame the discussion:

> *Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*

The reason that this notion of Cloud Computing is relevant for discussions about sustainability is that it allows users to access computing as part of a shared marketplace. By aggregating demand in central locations, Cloud Computing vendors can build data centers that use the best technology, located in the most sustainable locations and designed to achieve the greatest per-unit efficiency possible.

The following section will look at the details of these implications in detail.

The Positive Impacts of Consolidation

Consolidating disparate data centers into large scale shared premises logically reduces the net environmental impacts. But it is worthwhile to review the ways in which these benefits are realized in practice. In a recent report on the subject from Accenture, [5] broad categories were identified which drove these changes;

- Reducing total infrastructure allocation
- Leveraging multi tenancy
- Maximizing utilization rates
- Improving the data center efficiency

Let us look at these different areas in turn.

Reducing Total Infrastructure Allocation

Server, networking and storage usage is rarely stable. As a result, an organization's infrastructure deployment needs to be sufficiently large to allow for the maximum peak load that will be experienced and in many situations this can be many times the average load. This fact is compounded by the fact that capacity planning is very difficult and hence IT departments tend to allow themselves some "wiggle room" to ensure no service degradation or outages occur.

Unlike individual companies, however, Cloud providers are able to smooth the peaks of their various customers across their infrastructure. Since every organization has a different peak profile, that is their peaks occur at different times, it stands to reason that total infrastructure requirements will be less in an aggregate situation than when individually provisioned. This fact was highlighted by Joe Weinman in his 10 Laws of Cloudonomics [9]. Law three states that;

> *The peak of the sum is never greater than the sum of the peaks.*
>
> *Enterprises deploy capacity to handle their peak demands. Under this strategy, the total capacity deployed is the sum of these individual peaks. However, since clouds can reallocate resources across many enterprises with different peak periods, a cloud needs to deploy less capacity.*

Cloud providers therefore reduce the inefficiency created by individual data centers over provisioning by smoothing peaks across all users.

In addition to reducing total infrastructure allocation, one of the key drivers of efficiency gains in the cloud lies with multi tenancy, the next subject we discuss.

Leveraging Multi Tenancy

Multi tenancy is the approach towards infrastructure that sees multiple organizations use the same physical hardware. Virtualization enables this and is the technology that allows the;

> *...division of a single physical server into multiple "virtual" servers containing multiple sets of segregated data* [8]

Virtualization is absolutely fundamental to the efficiency gains from Cloud Computing. To illustrate this point, studies [10] have found that Cloud applications consume up to 90% less energy than on-premise ones.

A recent study commissioned by Microsoft to investigate the environmental impacts of Cloud applications [5], enumerated these benefits. While the study looked at application hosting specifically as opposed to infrastructure hosting, the benefits in terms of economic impacts that it found would apply to all areas of Cloud Computing. In the chart below the savings in terms of CO2 emissions gained from using three different Cloud applications are depicted for small, medium and large organizations. In some cases, the CO2 reductions were over 90% in a cloud environment.

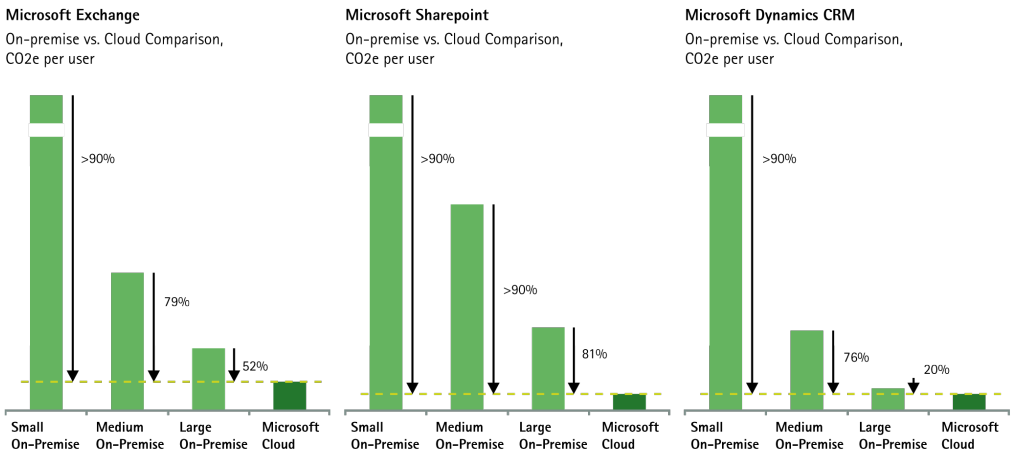

Comparison of Carbon Emissions of Cloud-Based vs. On-Premise Delivery of Three Microsoft Applications [11]

Related to virtualization is the third aspect of reduced environmental impacts, the fact that cloud providers tend to have better server utilization than standalone infrastructure.

Server Utilization

Cloud Computing providers utilize the aggregated demand over all their customers to run their servers at higher utilization rates, the percentage of computing capacity that is actually utilized on a server, than those in individual organizations. Because of this higher utilization, it requires fewer servers to perform the same tasks in a cloud environment than it does in a traditional setting. The following diagram [12] illustrates this fact. It should be noted that increasing the utilization rate from 5% to 15% allows a server to process many times the previous load yet increases power consumption by a relatively small amount.

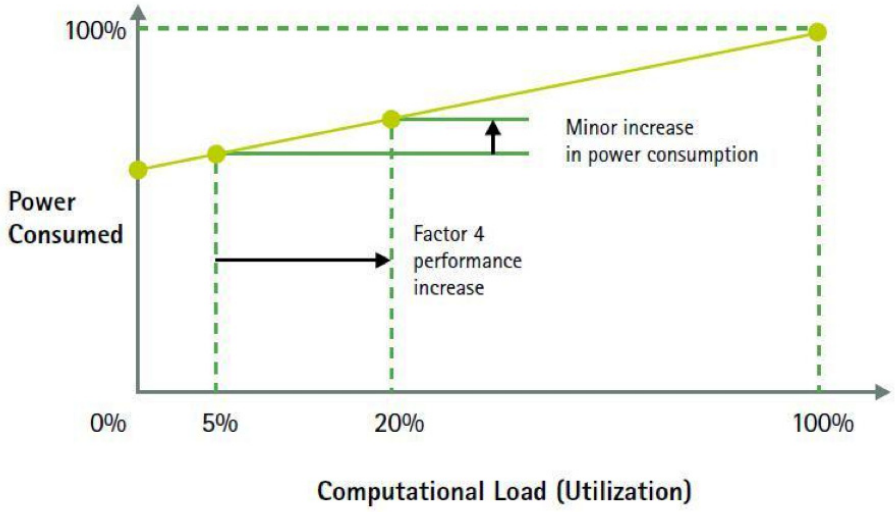

Relationship between Server Utilization and Power Consumption

By using a virtualized infrastructure, multiple instances of an application, or multiple servers, can run on the same piece of physical hardware. This machine can then be tuned to run at high utilization rates, thus reducing the total number of servers required to run a particular workload.

The last three factors for Cloud Computing sustainablity benefits related to the servers themselves; however Cloud Computing vendors also have an advantage when it comes to the overall design of the data center.

Data Center Design

Building and running data centers are two of the largest expenses that Cloud Computing providers incur. Making these data centers efficient is thus a top priority. PUE (Power Usage Effectiveness) is an efficiency benchmark comparing the power used to run a data center's infrastructure (Total Facility Power) to the power used to drive the IT equipment. [13] It is often used in discussions of sustainability and can be a useful metric but there are some caveats about PUE that should be considered.

PUE is a point in time measurement that reflects the current conditions in a data center. While it was not intended to use to compare one facility to another, PUE does help us understand what the power in a data center is being used for—essentially cooling or compute. To run a data center efficiently, more electricity dollars need to go toward compute and less to cooling and this is why PUE has received so much attention.

A PUE of 2.0 means that for every Kilowatt consumed, twice the amount of power is used to run the facility as is used to run the IT equipment. A PUE of 1 (an extreme example) would use identical amounts of electricity for both the facility infrastructure and the IT equipment.

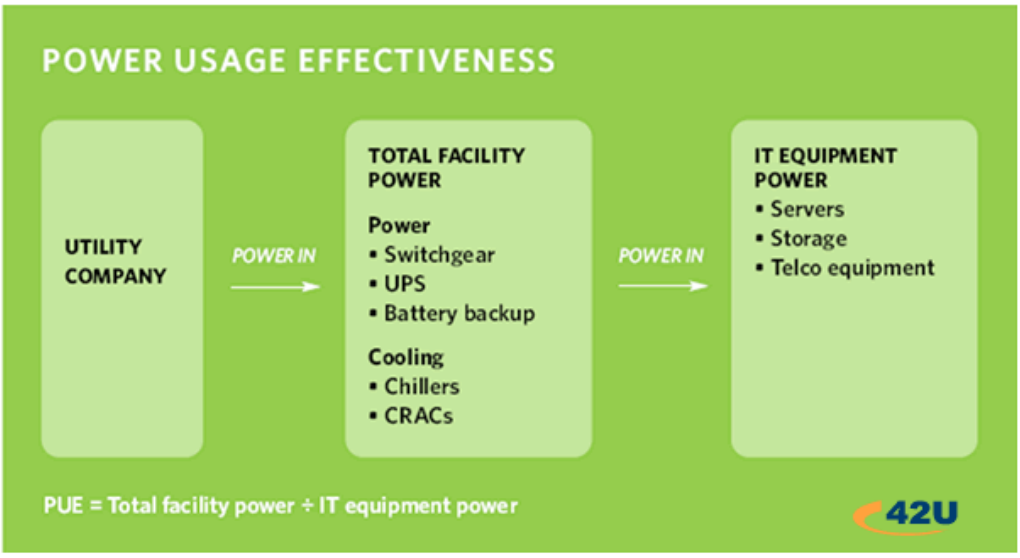

Power Usage Effectiveness - PUE (Source: Green Grid)

Since data centers consume electricity as a main resource, Cloud Computing providers have an incentive to design them to be highly energy efficient. Some of the technology required to drive these efficiencies can be impractical to install in all but the largest

facilities due to cost and complexity, giving Cloud Computing providers an advantage in maintaining cutting edge facilities. The US Environmental Protection Agency has reported [14] that average PUEs in 2009 for US data centers are around 1.97. Large scale Cloud Computing vendors aim for PUEs as low as 1.1 to 1.2 thereby significantly reducing per-unit impacts of IT processing and storage. To put that in context, a difference of 0.7 in efficiency for a 4 Megawatt load will drive an annual reduction in electricity of 28 Million Kilowatts, savings of $2.6 Million dollars and will avoid 16 tons of carbon.

While the case for more efficient data centers is compelling today, increases in the demand for electricity and the associated green house gasses generated to meet that demand will help drive the change that will align what people need with planet and profit.

A Changing Electricity Paradigm

Given the constraints of the current energy framework, there has been a trend towards building data centers where electricity costs are lower, though not always cleaner, or in places where climates are colder, using outside air for 'free' cooling eliminates the need for electricity to cool the facility. This trend shows that even today with the relatively low cost of electricity, increasing efficiency makes economic sense and drives business decisions.

We believe that the awareness to reduce carbon emissions will drive innovation to develop plentiful renewable sources of electricity within our generation. This will avert a scenario of ever increasing regulation and carbon taxes driving a more expensive electricity marketplace. The proactive development of less expensive, renewable sources will change the world for generations to come.

Summary

Economic predictions suggest that the price of energy will continue to rise. Simultaneously with this trend organizations will be looking to maximize efficiencies and minimize the environmental impacts of their activities.

With IT becoming an ever increasing driver of the economy, providers will look for ways to do more, with less. One of the ways to increase outputs with no increase in impacts is by moving computing away from individual data centers and into Cloud Computing.

Consolidated IT drive environmental benefits via;

- Reducing total infrastructure allocation
- Leveraging multi tenancy
- Maximizing server utilization rates
- Improving data center efficiency

As such we predict that the move to consolidated computing, especially Cloud Computing, will continue to increase in the years ahead.

About Diversity Analysis

Diversity Analysis is a broad spectrum consultancy specializing in SaaS, Cloud Computing and business strategy. Our research focuses on the trends in these areas with greater emphasis on technology, business strategies, mergers and acquisitions. The extensive experience of our analysts in the field and our closer interactions with both vendors and users of these technologies puts us in a unique position to understand their perspectives perfectly and, also, to offer our analysis to match their needs. Our analysts take a deep dive into the latest technological developments in the above mentioned areas. This, in turn, helps our clients stay ahead of the competition by taking advantage of these newer technologies and, also, by understanding any pitfalls they have to avoid.

Our Offerings: We offer both analysis and consultancy in the areas related to SaaS and Cloud Computing. Our focus is on technology, business strategy, mergers and acquisitions. Our methodology is structured as follows:

- Research Alerts
- Research Briefings
- Whitepapers
- Case Studies

We also participate in various conferences and are available for vendor briefings through Telephone and/or Voice Over IP.

About Rackspace

Rackspace Hosting is the world's leading specialist in hosting and Cloud Computing. The San Antonio-based company provides Fanatical Support® to its customers, across a portfolio of IT services, including Managed Hosting and Cloud Computing. Rackspace is also the founder of OpenStack™, an open source Cloud platform with broad industry support, designed to offer Cloud consumers greater choice. For more information, visit www.rackspace.com.

About the author

Ben Kepes

Ben is the founder and managing director of Diversity Limited, a consultancy specializing in Cloud Computing/SaaS, Collaboration, Business strategy and user-centric design. More information on Ben and Diversity Limited can be found at http://diversity.net.nz

1. http://www.earthday.org/about-us

2. http://www.ipcc.ch/

3. http://www.gartner.com/it/page.jsp?id=503867

4. http://www.gesi.org/

5. Cloud Computing and Sustainability: The Environmental Benefits of Moving to the Cloud, WSP, 2010

6. Nicholas Carr, *Big Switch: Rewiring the World, from Edison to Google*.

7. http://www.gartner.com/resources/208500/208540/sustainability_for_growth_a__208540.pdf

8. http://broadcast.rackspace.com/hosting_knowledge/whitepapers/Revolution_Not_Evolution-Whitepaper.pdf

9. http://gigaom.com/2008/09/07/the-10-laws-of-cloudonomics/

10. Cloud Computing Emissions Comparison, Nucleus Research, 2 Sept 2010

11. http://www.accenture.com/SiteCollectionDocuments/PDF/Accenture_Sustainability_Cloud_Computing_TheEnvironmentalBenefitsofMovingtotheCloud.pdf

12. http://www.thegreengrid.org/Global/Content/white-papers/Five-Ways-to-Save-Power; The Green Grid, Five Ways to Reduce Data Center Server Power Consumption, 2008

13. http://www.42u.com/measurement/pue-dcie.htm?source=what&#What-is-PUE

14. USEPA, Energy Star Data Center Rating, 2009