

Greedy for Green: The Motivations and Hindrances of Maintaining Environmental Sustainability in the Commercial Industry of Cloud Computing

Bradley S. Cantor
Enrolled in the Master of Software Engineering Program
Department of Computer Science
Villanova University
Villanova, Pennsylvania 19085, U. S. A.
bcanto01@villanova.edu

Abstract. This article explores the potential for environmental sustainability in the rapidly growing industry of Cloud Computing. By understanding the Cloud as a commercially driven endeavor, we unearth the motivations that drive the growth of the burgeoning industry. We go on to explain the impetus behind many of the Data Center innovations that Cloud Computing has brought about. We show that these innovations are why the Cloud has been heralded as a form of Sustainable Computing. After exploring the technologies endorsed and innovated by Cloud Computing, the article proceeds to discuss some of the potential downfalls inherent to propagating a commercially driven, yet environmentally sustainable industry. Lastly, the paper discusses the role that we, the consumer, play in shaping the environmental sustainability of the digital storage industry.

Keywords: Cloud Computing, Sustainable Computing, Data Center innovations, ethical consumerism, economic viability

1 Introduction

We live in an era of incredible innovation. The speed of this innovation is almost overwhelming. There are few among us who have not felt the remorse brought about from buying a brand new smartphone or flat-screen TV only to learn days later that a sleeker, fancier model has been announced. This type of rampant technological

growth defines the past few decades. Landlines have given way to cellphones. Dial-up modems are all but extinct. Wi-Fi flows wherever there is coffee (just be careful asking for 'Java'). High-schoolers can become 'friends' with a single click and a series of hash-tags is currently creating the next big, yet brief, internet trend. The speed of change is hard to keep up with and impossible to ignore. #innovation

Constantly evolving gadgets, software applications and websites have become an integral part of daily life. This has created a society in which technology is no longer optional. The mail cluttered kitchen counters of our parent's generation have been replaced by our own poorly managed and often overflowing digital inboxes. Our computers act as our rolodexes, our file cabinets, our calendars, book shelves and CD collections. The stuff that used to be jumbled around our homes now fills our hard drives. The movement to the digital environment has altered the way that we think about storing information. Fifteen years ago, nobody would have concerned themselves with bringing two thousand CD's and fifty seven books on a two hour train ride. The advent of smart phones and tablets has changed our perception of accessibility. We want everything we own to be everywhere we go, and we want to be able to access it instantly. The rigid upload/download format of yesteryears MP3 players has given way to a far more dynamic interface of personal data storage; the cloud.

Entire archives of data, artwork, and entertainment flow through a network of revolving doors that connect millions of computers, phones and handheld devices around the world. The sheer magnitude of data has turned 'space' into a digital commodity. It is sometimes easy to forget that this data does not just magically float around in the sky. The premium placed on space has created an entirely new storage industry. Data centers, or server farms, have begun to pop up all over the United States, Europe and Asia. These data centers power the cloud. Without them, there would be no Spotify, no Amazon Web Services, or Google App Engine. Every megabyte of information 'floating' in the Cloud is stored in one of these servers [1]. This means that as the Cloud continues to grow, more and more data centers will be designed, built, and utilized. This growth represents incredible commercial opportunity. We are already beginning to see a mad dash for tech companies to provide Cloud services. But the Cloud model embodies more than just economic potential. This new industry could very well help usher in a new age of corporate responsibility and environmental sustainability.

The way that data centers are designed, built, and utilized will determine both the economic and environmental potential of Cloud Computing. As the industry grows it has the capability to redefine standard best practices in the field of Sustainable Computing. Energy efficient data center design and innovative server technologies have seemed to evolve alongside the rapidly growing Cloud but these practices are in no way guaranteed. As with all commercial endeavors, profitability will inform practice. In the rampantly growing industry of digital storage, economic viability will supersede environmental ideology. The potential for the Cloud to usher in a new era of Sustainable Computing will rely not on ethical responsibility, but instead on dollars and cents. The question is not whether the ethicality of the Cloud can stand up to corporate consumerism but instead whether the two coincide.

2 Sustainable Computing

The field of Sustainable or ‘Green’ computing is “the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems—such as monitors, printers, storage devices, and networking and communications systems — efficiently and effectively with minimal or no impact on the environment[3]”. While this definition sums up many of the goals of Sustainable Computing there is another side of the environmental endeavor that often gets overlooked. Sam Murugesan, the author that provided the aforementioned definition, writes that “Green IT also strives to achieve economic viability and improved system performance and use, while abiding by our social and ethical responsibilities [3]”. This aspect of Sustainable Computing perfectly depicts what it is that makes the Cloud so interesting. The future of Sustainable Computing depends on the potential of economic viability. The mere practice of designing, using and disposing of hardware in an environmentally friendly way will not, in itself, provide an incentive for large-scale change. The same can be said for Data centers. If the green option is more expensive or less effective it will be unable to compete in the market. Cut-throat corporations are not going to trade economic viability for environmental sustainability. While this may seem problematic for green computing and the future of the Cloud, economic viability and environmental sustainability are not mutually exclusive.

3 Incentives for Cloud Computing

A survey conducted by Sun Microsystems Australia involving 1,500 responses from 758 large and small organizations in Australia and New Zealand, found that reducing power consumption and lowering costs were the major reasons that companies chose to incorporate environmentally responsible practices. The environmental impact of these practices was something of an afterthought [4]. Businesses prioritize environmental sustainability if and when it is cost effective. When corporate concerns (cost and profit) align with sustainable practices the motivation for these practices need not matter. Sustainable Computing on the corporate level does not need to be a morality contest. Corporations will choose eco-friendly practices when they are the most cost effective. That is the bottom line. The environmental benefits are merely a byproduct.

This is where the innovations of Cloud Computing come into play. The Cloud presents a potential for economic viability and Sustainable Computing to coincide. In the past, corporations, both large and small, relied on small on-site data centers. That meant that in some form or another, companies were forced to use office space to physically house their servers. The databases, websites, and software that allowed businesses to function were all stored in these localized servers. Any loss of server functionality was potentially devastating; an onsite crash could instantly debilitate a business. This meant that corporations had no choice but to heavily invest in the maintenance and upkeep of their onsite data centers. Every localized data center needed its own team of IT professionals to provide technical support for the servers. In addition to these IT professionals, these small in-house data centers required

housing specifications and environmental control systems that far exceeded the ordinary demands of generic office space [5]. Insulation and temperature control was vital to preserving optimal functionality. The introduction of Cloud Computing mitigates most, if not all, of these costs. Companies can spend less on IT professionals, optimize office space and, most of all, avoid incredibly inefficient energy expenditures.

By outsourcing digital storage, companies are able to streamline costs. According to a recent research report by Accenture, small businesses experienced a reduction in emissions of up to 90 percent while using Cloud resources. Additionally, large corporations also saw improvements of 30-60 percent in carbon emissions while using Cloud applications[6]. But why is the Cloud environmentally beneficial? It would seem that it just moves energy expenditure from one place to another. The Cloud model simply replaces two million, small-business onsite data centers with a few massive server farms. Perhaps there are fewer data centers but the amount of information that has to be stored is still the same. This would seem to suggest that the same amount of servers, and therefore, the same amount of net energy would still be required. After all, the sum of the parts is always equal to the whole. While it is hard to say that this line of thought is not highly logical, it is also flawed. The truth is that Cloud Computing represents far more than just a geographical relocation of data centers. Technological innovations brought about by the Cloud are changing the way that servers are stored and utilized. It is these innovations that are minimizing energy expenditure and therefore aligning corporate incentive with environmental ethicality.

4Innovations in Cloud Based Data Centers

Emerging standard best practices are enabling Cloud based data centers to become increasingly energy efficient. Pre-Cloud on-site data centers were designed to handle sporadic peak loads. The problem with this approach is that during non-peak hours there was no way to reduce operational functionality. This reduced resource utilization resulted in wasted energy. Data centers that use emerging cloud practices, on the other hand, can greatly increase resource allocation through server consolidation. Different companies can share the same server using a parallel processing and partitioning method called virtualization. This practice helps alleviate energy consumption that would have otherwise been spent on the electrical costs of running a local server during non-peak hours. As workloads are consolidated onto partitioned servers, unused servers can be switched off. Before virtualization, this approach was impossible. Each server in an onsite data center had its own data, its own files and therefore its own function. It is only the dynamic partitioning of virtualization that allows for the ebb and flow of server consolidation to conserve energy [7].

The innovations brought about by Computer Science are in themselves remarkable, but the Cloud also presents a potential for sustainability that is far more intuitive. Other than the power required to physically run the servers, the largest energy expenditure experienced by data centers is the cost of temperature control [8].

Before the Cloud, data centers were generally a converted room or collection of rooms in an office building. The center was almost certainly retrofitted to house servers; it was not the original purpose of the structure. The state-of-the-art server farms that are now being utilized by companies that provide Cloud services are cut from a very different cloth. Buildings are being designed for the sole purpose of storing servers. The energy costs of cooling the servers can be alleviated by choosing sites that remain cold throughout the year. Harnessing wind, shade, water, and ground temperature can help lessen operational costs. Due to the energy demands inherent to data centers, infrastructure siting is becoming increasingly dictated by climate and resource allocation. Investing in these state-of-the-art centers is not cheap but companies that provide Cloud services seem to be increasingly willing to take this approach; spend now to save later [9].

5 Concerns for Future Sustainability

Data-centers are the infrastructure of Cloud Computing. They play an integral and inevitable part in the business of digital storage. By taking climate and ecological factors into account, data centers are able to reach new levels of energy efficiency. This green infrastructure is not only cost effective, but it also provides a strong anchor for the future of Sustainable Computing. After all, these new data centers are the roots on which the Cloud must grow. But a recent Greenpeace report released in April of 2012 is far more skeptical about the inherent greenness of “the factories of the 21st century information age. [9]” Weary of the potential downfalls of massive data centers and the large corporations that own them, the Greenpeace report stresses the importance of energy transparency, infrastructure siting, and the use of clean/renewable energy. While the findings do suggest a trend towards corporate environmental advocacy, the report also argues that the Cloud phenomenon is not necessarily, in and of itself, eco-friendly.

It should come as no surprise that the issue is energy. The electrical cost of maintaining the aggregated digital information at present is already beginning to become astronomical. The expected growth of Cloud Computing over the coming years will dramatically increase this cost. According to the Environmental Protection Agency, “data centers now account for 1.5 percent of all electricity consumption in the U.S. and by 2020, carbon emissions will have quadrupled to 680 million tons per year, which will account for more than the aviation industry [10]”. Furthermore, a recent estimate of the IT sector’s footprint conducted as part of the 2008 SMART study concluded that data centers will be responsible for 2% of global GHG emissions by 2020 [11]. Incredibly, the Greenpeace report also estimates that the industry will experience “a 50-fold increase in the amount of digital information by 2020 and nearly half a trillion in investments in the coming year [11]”.

These numbers are staggering. They illustrate the profound technological and commercial role that Cloud computing will play in the coming years. In the face of such rapid expansion, Greenpeace, and many others, worry that Cloud providers will become increasingly fixated on lowering the operational energy costs of their data centers. There is no question that infrastructure siting and design innovations will help

accomplish this but in business, profit informs growth. State-of-the-art data centers are expensive and take years to begin to pay for themselves. In the mad dash to take up Cloud market shares there is a strong possibility that rising energy costs will be mitigated by purchasing cheaper fuel. Infrastructure siting is not just about harnessing climatic features; it is heavily influenced by geographical resources. A company that decides to build a new data center in West Virginia is doing so because of the price of coal, not because of the cool mountain air.

If coal becomes the main fuel source of data centers, the environmentally friendly practices innovated by Cloud technologies will be vastly overshadowed by the sheer quantity of greenhouse gases produced by the industry. Energy efficient techniques like workload consolidation and virtualization will not be enough to level out the environmental impact [12]. Fuel is the one fundamental issue in Cloud Computing in which economic viability does not align with sustainability. If data centers turn to coal, Cloud Computing will cease to be “ethically responsible”. The potential for a new age of Sustainable Computing brought about by the burgeoning Cloud will fall by the wayside. Simply put, a Cloud industry that garners the majority of its fuel from coal will be environmentally disastrous.

6 Consumer Ethicality

With the continual growth of the internet and the increasing market share of server-dependent companies, the demand for Cloud technologies and data centers is going to increase substantially. Cloud computing is in its embryonic stages. The possibilities are in many ways limitless, but as everyday consumers become increasingly accustomed to the possibilities of the Cloud, the temptation for data centers to tap cheap energy sources will prove difficult to resist. Low energy costs have the potential to subvert high energy efficiency. Cheap fuel will equal cheap digital storage. In an industry dictated by profit, it would seem that this business model would be difficult to combat. If offered the option to pay five cents or ten cents for a gig of space which would you choose? For the past few years Greenpeace has issued environmental report cards for Cloud Service Providers. Companies like Amazon and Twitter have literally failed in multiple categories. Yet, these are two of the most profitable companies in the Tech industry [11][13] [14].

As the Cloud continues to grow, economic viability will determine the policies and practices of service provider data centers. It is easy to say that large corporations should do more to protect the environment, but what is our role in all this? At the end of the day, it is not Google or Greenpeace that will decide the future of environmental sustainability in digital storage. It is the end user. It is whether or not we care. If the environment matters to the end user it will be economically advantageous to design, build, and utilize data centers in a way that coincides with the principles of Sustainable Computing. If the end user does not care, following these principles will cease to be economically pragmatic.

Companies like Yahoo, Facebook, and Google are spending millions of dollars on new, state-of-the-art, ultra energy-efficient data centers. By tapping renewable energy sources like nuclear, wind and solar, these companies are making

strides to help ensure that the Cloud remains a green enterprise [15]. Their actions set the bar high for others to follow, but these are companies that are never too far from the public eye. Appearing environmentally trustworthy is good for business. But branding, like energy conservation, helps align economic viability with ethical and environmental responsibility.

7 Conclusions

We use the Cloud to make our lives simpler—more streamlined, organized and accessible. Cloud Computing allows us to clean off our kitchen counters, throw out our scratched old CD's, and travel anywhere in the world with a library in our pocket. But rarely do we stop and think about the cost of this convenience. Rarely do we consider the equation in its entirety. What is the cost of this digital information?

Within the Cloud, every undeleted email, every blurry family picture, every video of a cat wearing socks, is stored on a server. Innovations in server technology along with data center design and utilization have helped forge a digital storage industry that is environmentally responsible and economically profitable. The Cloud has brought about innovation in infrastructure siting, workload consolidation, and server virtualization. But, the rapid growth of the industry has the potential to undermine the environmental benefits of these innovations. The Cloud is an energy hungry enterprise. If data centers turn to coal as the main source of energy, the Cloud will cease to be 'green'.

We as consumers hold the key to this green potential. The industry of Cloud Computing will become whatever we make it. Our decisions, our values, our beliefs, and, most of all, our dollars, will determine the future of the Cloud. The consumer will ultimately define the practices that are profitable. If cheap storage powered by coal is demanded, the market will be sure to supply it. While this understanding of the Cloud may lack a certain ideological flare for the ethics of capital markets, it in no way condemns the Cloud to a future of bleak possibilities. We as consumers should care about the environment, and perhaps we will. But the Cloud will not stay green because it ought to, or because Sustainable Computing is ethically correct. The Cloud will stay green if and only if it is economically viable.

Acknowledgments

I would like to thank Dr. William M. Fleischman for his support and guidance throughout this process. This article originated as a term paper in the Computer Ethics class offered through the Computer Science Department at the University of Villanova. Without Dr. Fleischman's foresight and curiosity, I would never have been able to develop that paper into the article that it has now become.

References

1. Anderson, S., Improving Data Center Efficiency, *Energy Engineering*, Vol. 107, No. 5, pg 42-63. (2010)
2. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., and Zaharia, M., A View of Cloud Computing, *Communications of the ACM*, April, Vol. 53, No. 4. (2010)
3. Murugesan, S.: Harnessing Green IT: Principles and Practices, *IEEE IT Professional*, January–February 2008, pp 24-33.
4. Andonova, L.: Networks, Club Goods, and Partnerships for Sustainability: The Green Power Market Development Group, Colby College
5. Greenberg, S., Mills, E., Tschudi, B., Rumsey, P., and Myatt, B., Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers. ACEEE Summer Study on Energy Efficiency in Buildings. available at <http://eetd.lbl.gov/emills/PUBS/PDF/ACEEE-datacenters.pdf>, last accessed 10 May 2013 (2008)
6. Accenture Microsoft Report: Cloud computing and Sustainability: The Environmental Benefits of Moving to the Cloud, available at http://www.wspenvironmental.com/media/docs/newsroom/Cloud_computing_and_Sustainability_-_Whitepaper_-_Nov_2010.pdf. last accessed 21 June 2013 (2010)
7. Verma, A., Koller, R., Useche, L., and Rangaswami, R., 2010. SRCMap: energyproportional storage using dynamic consolidation. 2010, Proceedings of the 8th USENIXconference on File and storage technologies (FAST'10), San Jose, California
8. Robles, L. Four Trends that Shaped Cloud Computing in 2011, available at <http://venturebeat.com/2011/11/29/four-trends-that-shaped-cloud-computing-in-2011/> last accessed 21 May 2013 (2011).
9. Greenpeace International: Make IT Green available at <http://www.greenpeace.org/international/en/publications/reports/make-it-green-Cloudcomputing/>, last accessed 21 July 2013 (2012)
10. Metreweli, K., Data Center Cost Savings Go Green, available at www.ebizq.net/topics/int_sbp/features/11740.html, last accessed 21 May 2013 (2009)
11. Greenpeace International: How Green is Your Cloud, available at, <http://www.greenpeace.org/usa/Global/international/publications/climate/2012/iCoal/HowCleanisYourCloud.pdf> last accessed 25 July 2013 (2012)

12. Garg, K., Buyya, R., Green Cloud Computing and Environmental Sustainability, The University of Melbourne, Australia
13. Forbes Magazine: Global 2000: The World's Biggest Public Companies available at <http://www.forbes.com/companies/amazon/>, last accessed 27 July 2013 (2012)
14. Shontell, A.: The Digital 100: The world's most valuable private tech companies, Buisness Insider 2013, available at <http://www.businessinsider.com/2012-digital-100?op=1>, last accessed 27 July, 2013 (2013)
15. Cook, G., Van Horn J. How dirty is your data?, available at <http://www.greenpeace.org/international/en/publications/reports/How-dirty-is-your-data/>, last accessed 21 July 2013 (2011)
16. Comerford, T.: Principles of Data Center Siting and Economic Development Incentives, UEDA Winter Forum, Feb 23, 2011 available at http://www.blsstrategies.com/Docs/Events/Event_33.pdf, last accessed 27 July 2013 (2011)