Environmental Modelling & Software 56 (2014) 6–12

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

The greenhouse gas abatement potential of enterprise cloud computing☆

Daniel R. Williams [a,*], Peter Thomond [b], Ian Mackenzie [c]

[a] Technologies for Sustainable Built Environments (TSBE) Centre, University of Reading, Reading RG6 6AF, UK
[b] Clever Together LLP, 5A Maltings Place, 169 Tower Bridge Road, London SE1 3JB, UK
[c] Harvard Business School, Soldiers Field, Boston MA 02163, USA

ARTICLE INFO

Article history:
Received 24 May 2013
Received in revised form
6 November 2013
Accepted 19 November 2013
Available online 9 January 2014

Keywords:
Environment
Cloud computing
GHG
Green ICT
Technology

ABSTRACT

In this paper we present a dynamic, country level methodology and model to determine the energy related Green House Gas (GHG) emissions abatement potential of enterprise cloud computing. The developed model focused upon demonstrating the impact of a move to cloud computing from traditional on-site computing, by creating country specific estimates of energy and GHG reductions. The methodology presented includes variables for market penetration, organisation size, and organisational adoption of on-site and cloud computing. Using the current enterprise cloud service applications of email, customer relationship management (CRM), and groupware against selected global countries, results indicated that 4.5 million tonnes of CO_2e could be reduced with an 80% market penetration for cloud computing. The majority of reductions were calculated to be from small and medium size organisations. A sensitivity analysis of the market penetration and current organisational adoption of cloud computing highlights the possible large variability in overall energy and GHG reductions. An analysis of the model and data used within this study illustrates a requirement for industry and academia to work closely in order to reach the large energy reductions possible with enterprise cloud computing.

1. Introduction

Cloud computing enables computing services (software, platforms, infrastructures, etc.), that are traditionally provisioned on-site within organisations, to be delivered on-demand from purpose built data centres across the internet. Past commercial studies (Accenture, 2011; Carbon Disclosure Project, 2011) have demonstrated how a shift from on-site to cloud computing can potentially reduce energy consumption and related GHG emissions. The main enabler of these reductions is the provision of computing resources from servers maintained in large centralised data centres; which, can be considerably more efficient than individual or small groups of non-virtualised servers, typical of on-site computing. Additionally, large centralised data centres can be located where energy costs and energy related GHG intensities are lower. Therefore, a shift to cloud computing, from on-site computing, could deliver significant reductions in energy related GHG emissions.

Previous studies (Accenture, 2011; Carbon Disclosure Project, 2011) have made a valuable start in quantifying the GHG abatement potential of some cloud computing services, calculating an abatement potential of between 50% and 90%. Yet, they leave many unclear assumptions and lack scientific data. In particular, little consideration has been placed towards how energy and GHG reductions are likely to vary by country or organisation size. Additionally, the studies do not publicly present scientifically appraised analysis methods, and fail to fully disclose market penetration assumptions for cloud computing. These issues limit the scientific usefulness for advancing the created methodologies and models, and have created the requirement for an open scientific analysis of the topic.

The goal of this research is to understand the GHG abatement potential of enterprise cloud computing when moving from an on-site technology base. Therefore, a methodology and model will be developed to calculate the energy consumption of computing services provided by on-site servers and cloud computing servers. Crucially, a variable to reflect the adoption of cloud computing will be included, which will allow a move from on-site to cloud computing to be assessed according to market penetration. The developed methodology will include variables to ensure that the

☆ Thematic issue on Modelling and evaluating the sustainability of smart solutions.

* Corresponding author. Tel.: +44 7736 060 870.
E-mail addresses: dw10@hotmail.com (D.R. Williams), peter.thomond@clevertogether.com (P. Thomond).

1364-8152/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.envsoft.2013.11.012

results can be tailored towards specific countries. Also, the design of the methodology will take into consideration that the efficiency of on-site computing varies by the size of computing resource, which in turn varies by the size of the organisation using the computing resource. This consideration will highlight if the proportion of different organisation sizes, which varies by country, makes a significant impact. Furthermore, the GHG intensity of energy provision varies by country as some have relatively high GHG energy (e.g. where coal fired power stations are significant) while others have relatively low GHG energy. Hence, the developed methodology will include variables for the profile of organisation sizes, their proportions within the countries being analysed and the relevant energy mix of each, thus providing more accurate country specific result.

A model will be developed from the methodology to scientifically demonstrate the value of a move to cloud computing, from on-site computing, by creating country specific estimates of energy and GHG reductions. The results of this model will be aimed at informing decision makers including ICT vendors, cloud service users, and governments alike. The model will also be used to perform a sensitivity analysis of the results to changes in key variables. The results of this will enable cloud vendors to ensure their cloud services are focused upon variables that maximise carbon abatement.

2. Methodology

An industry derived methodology developed by the Global e-Sustainability Initiative (GeSI) (GeSI, 2011) was utilised to provide a basic foundation and methodological framework to the study. The basis of the GeSI method is to assess the energy and GHG impact of information and communications technologies in the following three steps, which were followed in this research;

1. Define the aim and scope of the technology being studied;
2. Focus the assessment:
   a. Define a business as usual (BAU) baseline from which the new technology can be analysed from;
   b. Identify and focus upon the sections of the technology process that significantly impact energy consumption and GHG emissions;
3. Rigorously model and assess the significant changes that would be caused through the replacement of the BAU technology, and carefully interpret the results.

2.1. Scope

The goal of this research was to calculate the GHG abatement potential of enterprise cloud computing. Enterprise cloud computing was focused upon as it was a technology that could directly replace a BAU technology, which could be measured in the form of on-site computing. Also, enterprise level computing accounts for a large majority of computing usage, and richer data sets are available. In order to reduce a potentially huge analysis scope, the model focused upon three common applications that can be used in either on-site or cloud computing: email, customer relationship management (CRM), and groupware. While these are clearly a sub-set of all applications that organisations and individuals use, they are readily available as cloud based services and believed to represent a likely first phase in a shift towards cloud computing. Indeed, many more enterprise cloud service applications exist currently, and many more are likely to be available in the future.

The countries included in the model were Canada, China, Brazil, the Czech Republic, France, Germany, Indonesia, Poland, Portugal, Sweden and the United Kingdom (UK). These countries were selected as together they represented a broad mix of economies, and their energy supply GHG intensities varied. Baseline country data was collected for the year 2009 and 2007, dependent upon availability. More recent data for all countries and variables was not available at the time the analysis was carried out.

For the model, a cloud computing adoption rate of 80% was used for all countries. This ‘high’ value was selected to highlight a ‘best case scenario’ (Rogers, 2003), and one that is likely to take some years to achieve. A sensitivity analysis of this rate was also completed in order to assess its impact on the level of GHG abatement potential. The current or baseline adoption of cloud computing was also included for each country and organisation, so as to exclude impact already created by early adopters.

To make the inclusion of organisation size tractable, the size profile of organisations in each country was simplified by splitting them between small, medium, and large (s, m, l), drawing the boundaries from national statistics offices (Eurostat, 2010; LABORSTA, 2007; Statistics Canada, 2007; US Census Bureau, 2007). Furthermore, only organisations where most people actively use internet connected computers during their normal work routine were included in the analysis (defined in the data section).

2.2. Focussing the assessment

This research focuses upon the data centre stage of the overall cloud service data route. Cloud and traditional on-site networked data services follow a logical three stage data process route involving a server (within a data centre), network transmission, and end user devices (e.g. laptops, mobile devices). However, as established in previous research (Williams and Tang, 2012), the data centre stage of a digital process can account for the bulk of the energy demand compared to the network or user device stage. Whilst the energy consumption of the communications networks will likely be enhanced under a cloud computing scenario, preliminary estimates indicate that energy consumption within the network for the purposes in our scope would not materially impact overall energy consumption because of the small data sizes involved ((Baliga et al., 2011; Williams and Tang, 2013)). Finally, we assumed the populations and demands for end use devices would remain relatively unchanged by a shift to the three cloud based services (initially at least).

Cloud computing offers servers that are utilised more efficiently than standard on-site servers (Baliga et al., 2011; Berl et al., 2010; Kumar and Yung-Hsiang, 2010), and as such an assumption was made when calculating model results that on-site servers are switched off in favour of cloud bases services. Further to this, our analysis focused on the energy consumption involved in running servers, specifically the use phase. The embedded GHG content of the servers, the associated server maintenance and the energy implications of end-of-life decommissioning were all deemed to have small impacts relative to direct energy consumption (Bottner, 2010). Nevertheless, the embedded GHG content of servers included in (new) cloud facilities was estimated through a life cycle approach.

2.3. Modelling energy consumption

A Microsoft Excel based model was created to calculate the annual GHG emissions of on-site computing, and the GHG emissions of the equivalent services in a cloud computing scenario (Fig. 1) for a specific county, according to the defined scope and focus of the research. In order to begin the development of the new model, previous academic studies (Baliga et al., 2011; Berl et al., 2010; Kumar and Yung-Hsiang, 2010; Williams and Tang, 2013) were collated and interrogated for use in this model. These studies offered essential model variables to include, however, all were focused upon the detail of the cloud service instead of an organisational or national level view.

An example of the following calculations can be found in the supporting documentation.

2.3.1. Energy consumption for on-site servers

The annual energy (kWh) consumption for all on-site servers according to each on-site service (E_{ops}) was calculated using Eq. (1).

$$E_{ops} = (OS_{ops} * O_{ops}) * AE_{ops} \tag{1}$$

Where OS_{ops} is the number of servers per organisation, O_{ops} is the number of organisations using an on-site server, and AE_{ops} is the annual energy consumption of the on-site server.

In order to understand the scale of on-site server power consumption, the number of servers for each organisation size (s, m, and l) was required. Data on the exact number of servers for each organisation across each country was not available to the research. Therefore, the average number of servers per organisation (OS_{ops}) was estimated using Eq. (2).

$$OS_{ops} = C_{ops}^{-1} * (\overline{\text{emp}} * \text{ead}_s) + 1 \tag{2}$$

Where C_{ops} is the average session capacity of a server, $\overline{\text{emp}}$ is the average number of employees per organisation, and ead_s is the employee adoption rate for each service. The session capacity of a server reflects the average amount of user sessions a server can host at one time. This varies by the size of each user session and by the type of service being provided. The average number of employees per organisation and the employee adoption rate were required to determine how many servers were required for each organisation size against the session capacity. The employee adoption rate factored in the employee use of each service (S) per organisation, whether for cloud or on-site computing.

The number of organisations using an on-site server (O_{ops}) was determined in order to reflect the realistic use of on-site servers for a population of organisations. This was calculated using Eq. (3).

$$O_{ops} = \left(\frac{IW * \text{empshare}}{\overline{\text{emp}}}\right) * \left(\text{oad}_s * \text{opr}_{ops}\right) \tag{3}$$

Where IW is the number of internet workers, empshare is the proportion of total employees that work in either s, m, or l organisations, $\overline{\text{emp}}$ is the average number of