

2011

Furthering the Growth of Cloud Computing by Providing Privacy as a Service

David S. Allison

Western University, dallison.uwo@gmail.com

Miriam AM Capretz

Western University, mcapretz@uwo.ca

Follow this and additional works at: <https://ir.lib.uwo.ca/electricalpub>



Part of the [Computer Engineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

Citation of this paper:

Allison, David S. and Capretz, Miriam AM, "Furthering the Growth of Cloud Computing by Providing Privacy as a Service" (2011). *Electrical and Computer Engineering Publications*. 39.
<https://ir.lib.uwo.ca/electricalpub/39>

Furthering the Growth of Cloud Computing by Providing Privacy as a Service

David S. Allison, Miriam A. M. Capretz

The University of Western Ontario
London, ON, Canada
{dallison, mcapretz}@uwo.ca

Abstract. The evolution of Cloud Computing as a viable business solution for providing hardware and software has created many security concerns. Among these security concerns, privacy is often overlooked. If Cloud Computing is to continue its growth, this privacy concern will need to be addressed. In this work we discuss the current growth of Cloud Computing and the impact the public sector and privacy can have in furthering this growth. To begin to provide privacy protection for Cloud Computing, we introduce privacy constraints that outline privacy preferences. We propose the expansion of Cloud Service Level Agreements (SLAs) to include these privacy constraints as Quality of Service (QoS) levels. This privacy QoS must be agreed upon along with the rest of the QoS terms within the SLA by the Cloud consumer and provider. Finally, we introduce Privacy as a Service (PaaS) to monitor the agreement and provide enforcement if necessary.

Keywords: Cloud Computing, Privacy, Quality of Service, Privacy as a Service

1 Introduction

Cloud Computing represents an evolution of both computer hardware and software, as businesses and individuals alike no longer need to design, purchase, setup or maintain their own systems. Both hardware and software can be virtually maintained on the Cloud by a provider. Cloud Computing is desirable due to its cost effectiveness; the computing resources Cloud Computing can provide are often offered as a pay-as-you-go plan. In order for new computing evolutions, such as Cloud Computing, to gain widespread acceptance, the concerns consumers have in the technology must be addressed. Cloud Computing has shown success with consumers in some areas, such as delivering Web based email and online documents. These results, as seen in Google's Gmail [1] and Google Docs [2], are just a few of the possible uses for Cloud Computing. In order for Cloud Computing to continue its initial success and gain more widespread acceptance, major areas of consumer

Originally published in:

Information and Communication on Technology for the Fight against Global Warming, Lecture Notes in Computer Science Volume 6868, 2011, pp 64-78.

The final publication is available at:

http://link.springer.com/chapter/10.1007%2F978-3-642-23447-7_7

concern must be found and addressed. A 2010 survey of 100 IT professionals found that security was the top Cloud related concern, cited by 73% of the respondents [3]. Security is a far reaching topic in Cloud Computing and e-services, and includes such topics as authentication, authorization, auditing and privacy.

Privacy is a difficult topic with many unique problems. Privacy is subjective, as what can be considered private is unique to each individual. Providers prefer little consumer privacy, as the more information about a person a provider knows, the better it can create direct advertising. Due to these problems, of all the different topics in security, privacy is the least addressed [4]. Privacy is a particular concern to Cloud Computing, as Cloud providers necessarily have access to all of a consumer's data, and can use or disclose that information for unauthorized purposes, either accidentally or deliberately [5]. By addressing the issues of privacy, Cloud Computing will further gain the trust of consumers. This increase of consumer trust will give Cloud Computing a wider acceptance and will lead to its further growth as a technology.

It is important for the adoption of Cloud Computing to increase not only for economic reasons, but environmental as well. Cloud Computing advocates the better management of resources, resulting in the reduction of carbon emissions and the environmental impact of IT [6]. The environment impact of IT is substantial, accounting for 2% of all global carbon emissions [7]. A 2010 study commissioned by Microsoft [8] and conducted by Accenture [9] and WSP Environment & Energy [10], found that moving business applications from in-house to the Cloud can save a substantial percentage of carbon emissions, depending on the size of the business. Smaller businesses show the greatest benefit, reducing carbon emissions by up to 90 percent. Medium businesses are able to produce a 60 to 90 percent reduction, and large businesses a 30 to 60 percent reduction [11]. These reductions in carbon emissions were in large part thanks to four key aspects of Cloud Computing: the reduction of over-allocation of infrastructure, the sharing of application instances between consumers to reduce peak loads, the increased utilization of server infrastructure, and the improved efficiency of data centers to reduce the power required for cooling and maintenance [11].

In this work the ideas and issues surrounding privacy in Cloud Computing will be discussed, and the beginning of a framework to address Cloud Computing privacy will be presented. The goal of this work is to increase the adoption of Cloud Computing by providing privacy protection, in order to increase the economic and environment benefits Cloud Computing provides. We have previous experience in developing a privacy framework containing a Privacy Service and privacy policies in the related field of Service-Oriented Architecture (SOA) [12]. This work has the additional goal of being the first step towards a larger privacy framework for Cloud Computing. Section 2 discusses the current growth of Cloud Computing, and how the public sector and privacy can increase this growth. Section 3 provides an introduction to privacy, by first defining privacy, and then defining the different classifications of private information. Section 4 extends current SLAs with privacy quality of service (QoS) parameters. To this end, privacy constraints are developed from general privacy guidelines. These privacy constraints are defined, and then introduced into the SLA as QoS level parameters. Section 5 discusses the need for monitoring and legislation related to privacy in the Cloud. It is in section 5 that the concept of Privacy




Fig. 1. Services available in the Cloud

as a Service (PaaS) is introduced. This PaaS adds an extra level to the Cloud, which monitors the agreed upon privacy terms of the SLA. If any infractions are detected, the PaaS will provide alerts and enforcement if necessary. Section 6 discusses related work in the field of Cloud Computing privacy. Finally, Section 7 presents conclusions and future work.

2 Cloud Computing Growth

Cloud Computing has shown much growth since its inception less than a decade ago. However there is still much room available for continued expansion. In this section the abilities of Cloud Computing that have lead to its initial success are discussed, along with the Cloud's potential future growth and the role the public sector needs to play in that growth.

2.1 Current Reasons for Growth

The main advantage of Cloud Computing is its ability to provide an infinite amount of computer resources on demand, scaling to fit each individual need [13]. This relieves the Cloud Computing consumer from the task of planning ahead for future hardware and software requirements. Each Cloud consumer is charged for only what they are using, creating a renting structure for computer resources. Cloud Computing has been successful since its inception due to the large variety of services it provides. These services, as shown in Figure 1, range from low to high complexity and allow for a wide assortment of solutions to consumer needs.

- *Software as a Service (SaaS)*: SaaS allows a consumer to access the functionality of a software application over the Internet or any computer network. SaaS has the benefits of not requiring a consumer to install any software on their own computer, allowing the application to be accessed easily from multiple devices. SaaS is also economically friendly, as consumers do

not have to purchase an expensive software license, but rather lower priced access bundles based on access time or number of uses [14].

- *Infrastructure as a Service (IaaS)*: IaaS allows a consumer to gain access to a bare but complete hardware computing package. This includes hardware (HaaS), data storage (DaaS), network connectivity, and some fundamental software (SaaS), such as an operating system [15].
- *Hardware as a Service (HaaS)*: HaaS is made available through the virtualization of computer hardware over a network [16]. Access to physical hardware is made available to the Cloud consumer, who is charged based on the amount of processing power they utilize. Hardware virtualization allows the provider to use as many pieces of physical hardware as required to satisfy the demand of the consumer, while from the perspective of the consumer, only one piece of hardware is being used. HaaS provides a consumer with hardware flexibility and scalability not possible with on-premise hardware. The consumer always has access to the correct amount of hardware to meet their needs [16].
- *Data as a Service (DaaS)*: DaaS, also known as Storage as a Service, allows consumers to send and retrieve their data on externally provided storage. DaaS is uniquely defined as providing a consumer with the ability to create, read, update and delete data, rather than providing computation on data [17]. Through the use of DaaS, consumers have access to a virtual, scalable hard drive that will theoretically never run out of disk space. From the perspective of the consumer, their data is available just as if it were available on a local disk [16].
- *Platform as a Service (PaaS)*: PaaS is a remote computing environment delivered by Cloud Computing through the combination of HaaS, DaaS and SaaS [16]. PaaS builds on top of IaaS, by adding a layer of abstraction to automate the system, typically an application environment [15]. Consumers can subscribe to PaaS and gain access to a virtual platform for application development and deployment [18].

2.2 Future Growth and the Public Sector

The ability of the Cloud to provide remote, scalable, on-demand services at varying levels of complexity has led to its initial success and growth. This growth has been so strong, that some leading technologists have predicted that within the decade upwards of 90% of the world's computing and data storage will take place in the Cloud [19]. While it has become easier to predict how Cloud Computing will be used in the future, there are still many unanswered questions about the Cloud's future legal, economic and security details [19]. Privacy is one large aspect of security that needs to be addressed. In order to achieve a privacy solution, help from the public sector will be required. The public sector can adopt not only Cloud Computing, but standards and laws for providing privacy in the Cloud as well. This adoption by the public sector will also add to consumer trust of Cloud Computing technologies, leading to its further growth.

The public sector's involvement in Cloud Computing can produce a significant impact on the pace of the technology's development. Governments have shown the

ability to provide this impact in examples such as the adoption of telecommunication standards and the investment in needed infrastructure. As well, governments are often the largest economic entity in their country, and provide an example for other sectors of business [19]. For another example of the impact governments and public sectors can have on the growth of a new technology, one can look at the rapid growth and success of the Internet. Through the Internet's roots as the research of the United States Department of Defense, to the United States federal government's early adoption of Web sites and Internet Protocol [19], the public sector of the United States fueled the widespread growth of the Internet. Similar interest from public sectors around the world in Cloud Computing along with privacy protection, will greatly increase the technology's growth and adoption.

3 Privacy

In order to address privacy in Cloud Computing environments, one must first address the problem of defining privacy. Privacy as a concept is subjective, and has no single definition. In this section a definition for privacy will be given, which defines how privacy is viewed in this work. Also, the different types of private information one can expose on the Cloud will be defined.

3.1 Definition of Privacy

Security in computing is a large subject consisting of many different areas, such as authentication, authorization, auditing and privacy. Of these security concerns, privacy is the most difficult to define. Authentication is the process of determining that someone is who they claim to be. Authorization is the act of determining that a person or thing has the right to access a particular resource. Auditing is the task of recording the actions of a person or thing, ensuring that this person or thing cannot perform an act, and then later claim that it did not happen. These definitions are simple and unambiguous. Privacy is unique in that it has no one definition, the idea of privacy has changed and evolved over time. In 1888, privacy was defined as "the right to be left alone" by Justice Thomas M. Cooley [20]. This definition evolved into the ability to control the release of information about oneself. In the modern era, new technologies such as Cloud Computing have made releasing information about oneself often not an option, but a necessity of communication. In many cases this release of information is done without an individual's knowledge or consent. For this reason, in this work privacy is defined as the ability to protect information about oneself, and to also have some level of control over any information that has already been released.

3.2 Private Information

It is impossible to universally identify all information that should be considered private. What one person considers private may not be considered private by another. For example, some people freely list their telephone number in directories, while

others withhold this information. In lieu of a list of all forms of private information, it is important to identify different classifications of private information. As privacy is subjective and often relies on a given context, the classification definitions given in this section are not meant to be universal, but to apply in context to this work.

- *Personally Identifiable Information (PII)*: Personally identifiable information is any single piece or combination of information that can uniquely be traced to an individual. Some examples of a single-piece PII include credit card numbers, social insurance numbers, license plate numbers and fingerprints. Combination type PII include any grouping of information that together are associated with a single person. For example a name and a birthday individually may each point to several people, while together they can be used to find a specific person.
- *Sensitive Information*: Sensitive information is a classification that can be associated with a large number of people, but is still considered private by many due to personal concerns or personal preference. Examples of sensitive information include wage, age, sex, religion, and sexual preference. It is important to note that sensitive information can become PII if used in combination, or even by itself if the sample size is small or not diverse. For example in a city of people, knowing the age of a person would not be enough to identify that person. However in a workplace, there may only be a single person of a given age and thus age becomes PII.
- *Usage Information*: Usage information is gathered by tracking the history of any activity of an individual. Today, this is most commonly done through the tracking of an individual's activities on the Internet. When collected in small amounts, usage information cannot be traced to an individual. However when collected in large amounts over a period of time, even usage information can be used to deduce the identify of a person [21], thus it can become PII.

4 Extending the SLA with Privacy

The first step that is required to provide privacy in a Cloud Computing environment is a formal agreement between the Cloud consumer and provider, outlining how private information of the consumer will be handled by the provider. The privacy specifications of the Cloud consumer and Cloud provider will often be different and in conflict with one another. Therefore a negotiation process will need to be completed between the two parties. Once negotiations have been completed, the agreed upon privacy details will form a privacy contract. Cloud Computing currently contains a contract known as a Service Level Agreement (SLA). The SLA details the quality of service agreements between a Cloud consumer and provider, outlining the conditions under which a service can be provided to a consumer [22]. Typically, these SLAs describe technical details such as availability, accessibility, throughput, and response time. SLAs rarely, if ever, discuss privacy. Thus we propose the current standard of Cloud Computing SLAs should be expanded to include privacy terms and conditions.

4.1 Developing Privacy Constraints

In order to determine what privacy constraints should be added to the SLA, it is important to look at the current state of privacy legislation. The Organisation for Economic Co-operation and Development (OECD) has developed a set of Fair Information Practices (FIP) [23] which have been used as the basis for most of the privacy legislation throughout the world [24]. The FIP outlined by the OECD are a set of standards that govern the issues of privacy, both for the gathering and usage of personal information. The OECD guidelines produce eight basic privacy principles:

1. *Collection Limitation Principle*: Limits should be placed on the collection of any personal data. Data that is collected should be gathered legally with the knowledge or permission of the data subject.
2. *Data Quality Principle*: Personal data can only be collected if it is relevant to the purposes for which it is required. The collected data must also be current, whole and accurate.
3. *Purpose Specification Principle*: Any and all purposes for which the personal data is being collected should be specified at or before the time of collection. Any future changes of purpose must also be reported to the data subject.
4. *Use Limitation Principle*: Any personal data that is collected will not be made known or used for any purposes other than those specified by the Purpose Specification Principle. Exceptions to this are if consent is given by the data subject, or if the request is made with the authority of law.
5. *Security Safeguards Principle*: Collected personal data must be protected against dangers by reasonable security safeguards. Examples of dangers are: unauthorized access, alteration, removal, use and data leaks.
6. *Openness Principle*: Data controllers should provide transparency to their data collection by providing information regarding any data related practices, policies or developments. Data subjects should be provided the means to inquire about the existence of any personal data, the type of data, the purpose of use, the identity of the data collector and the location of the data.
7. *Individual Participation Principle*: Data subjects should be able to determine if any of their information has been gathered by a data controller. If any data has been collected, the data subject can request to be sent the data in an understandable format and in a reasonable amount of time. If the controller refuses either of these requests, the decision must be communicated to the subject and be challengeable. Data subjects should be able to challenge the information that has been gathered about them and if proven correct, have that data changed, removed or amended.
8. *Accountability Principle*: Accountability should be present to ensure the data controller is fulfilling all the above principles.

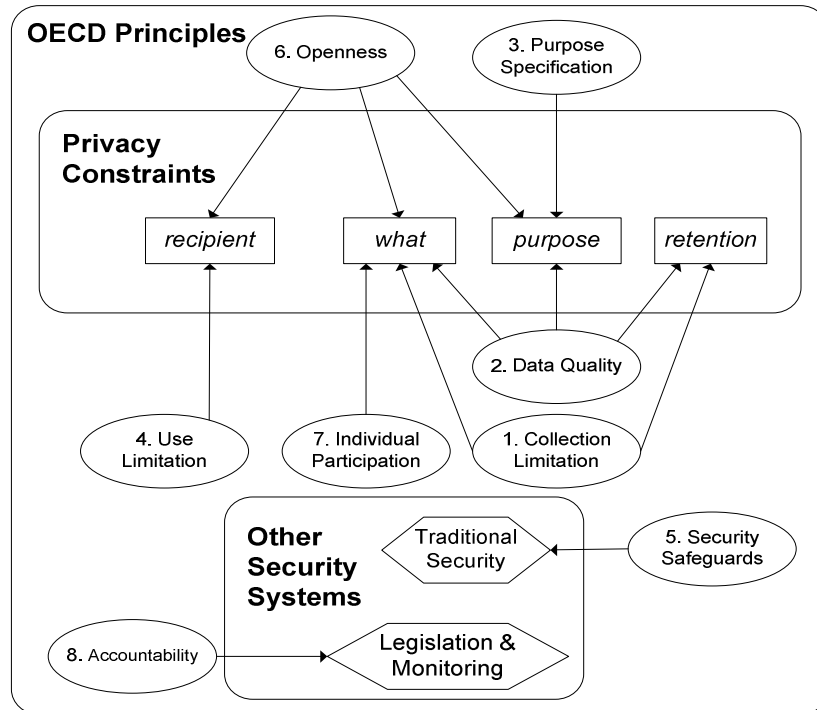


Fig. 2. Extracting privacy constraints from OECD principles

It is from these eight principles that the information that should be negotiated as part of a Cloud Computing SLA can be gathered. The process of selecting privacy constraints from these eight principles is shown in Figure 2 and is described in detail below. We have also used the same OECD guidelines in previous work to develop similar privacy policies for SOA [25].

- *Recipient:* The Use Limitation and Openness principles require the identity of the party who is allowed access to the private data be specified. This allows the proper party to not only gain access to the data, but also be available for further questions and challenges related to their data collection. Similarly, the Use Limitation principle states that there should be limits to whom the data can be disclosed. While the Openness principle requires the data controller be able to state the identities of all parties who have access to the data. From these two principles it was determined that the recipients of the data must be disclosed. This could be the single provider alone, or include parties the provider may pass information to. The recipient constraint allows this information to be known.
- *What:* Not surprisingly, the most common theme throughout the OECD guidelines is dealing with the data itself. The Collection Limitation principle expresses that the data subject must know what parts of their information are being collected. The Data Quality principle states that the data must be complete and accurate, while the Openness principle dictates that the nature of

the data must be made available. Finally, the Individual Participation principle lists challenges the data subject should be allowed to make in regards to their own private data. From these principles it becomes clear that what type of data will be collected must be defined. Only after this has been agreed upon between Cloud consumer and provider will these privacy principles be satisfied.

- *Retention*: Another requirement seen in multiple privacy principles deals with the idea of time. Collection Limitation states that there should be limits placed on the data collection, time being one such limit. Similarly, in order to keep the data up-to-date, as specified in the Data Quality principle, the age of the data must be specified. An agreed upon retention time would allow the appropriate length of time for storage of the collected data to be specified.
- *Purpose*: The Data Quality, Purpose Specification and Openness principles all require the reasons for which the data is collected be detailed. By outlining a purpose for the data collection, it can be assured that the possible uses of the data are known to both the Cloud consumer and provider.

Not every principle outlined in the OECD guidelines has been addressed by the requirements outlined above. This is because a privacy solution can only fulfill every privacy concern when included within a larger security framework. The Security Safeguards principle states that the data must be protected against unauthorized access and release. These concerns are addressed through the use of traditional security techniques, such as authentication, authorization and encryption.

The Accountability principle states the more abstract concern of holding the Cloud provider responsible for complying with all the other principles. Accountability presents a unique problem for Cloud Computing as the ability to provide enforcement is difficult and often nonexistent. This problem must be addressed through the combination of effective legislation and a Cloud monitoring system.

4.2 Privacy Quality of Service Levels

Cloud Computing environments implement SLAs in order to control the delivery and use of computing resources from a Cloud provider to a Cloud consumer. An SLA is defined by a schema containing Quality of Service (QoS) parameters. We propose the expansion of traditional SLA schemas to include a QoS privacy parameter. This privacy parameter will consist of four ordered levels of service. These levels can be organized by the provider to meet their needs. An example set of QoS levels is shown in Table 1. Each privacy constraint as defined in section 4.1 has a different value for each level of service. This expansion of the SLA with QoS privacy parameters was inspired by our previous work in creating metadata for Quality of Security Service (QoSS) for SOA [26].

The constraints in each QoS level are flexible, and can be changed to meet the requirements of any Cloud Computing environment. The constraints in each QoS level are also expandable and can be further defined. For example, the Cloud consumer and provider can create a further definition for the Data Category

Table 1. Example QoS privacy levels

Level	Recipient	Data Category	Purpose	Retention Time
High	Local	Consumer Specified	No Collection & No Distribution	7 days
Moderate	Trusted	Consumer Specified	Collection & No Distribution	30 days
Low	Enterprise	Provider Specified	Collection & Limited Distribution	365 days
Guest	Anyone	Not Specified	Collection & Distribution	Indefinitely

constraint. The data could be classified into the three categories of personal data: PII, Sensitive Information, and Usage Information. Each of these categories can contain whatever type of information is desired, and this determination is dependent on the perspective of the consumer and provider. With this extension, the consumer and provider can outline fine-grained privacy for specific types of data.

To help illustrate the example shown in Table 1, the Cloud consumer could select the moderate level of service from the provider. This selection would mean that the consumer allows the recipient to be anyone trusted by the consumer, meaning any other service the provider includes in the Cloud. The data category is consumer specified, meaning the Cloud consumer chooses what types of data they will allow to be collected. The purpose is for collection and no distribution, meaning the provider can read their data, but not share it with anyone. Finally the retention time outlines that the data can be held for a maximum of 30 days.

The QoS privacy parameters should be added to the schema in the same format as the current schema implementation. With the addition of the new privacy constraints, when the SLA is negotiated between consumer and provider, privacy will now be considered. The sophisticated process of SLA management and negotiation [6] is outside the scope of this work, but since the privacy QoS data is added to the schema in the same format as the preexisting SLA conditions, whichever negotiation process is currently applied should be easily convertible to handle the new constraint.

5 Monitoring and Legislation

While the negotiation of privacy terms within an SLA is necessary, it is not enough to provide adequate privacy in Cloud Computing environments. There must be a system in place to both monitor the status of the SLA, and to provide enforcement of its terms. In this section, we discuss how this monitoring can be done, and how legislation is required for enforcement.

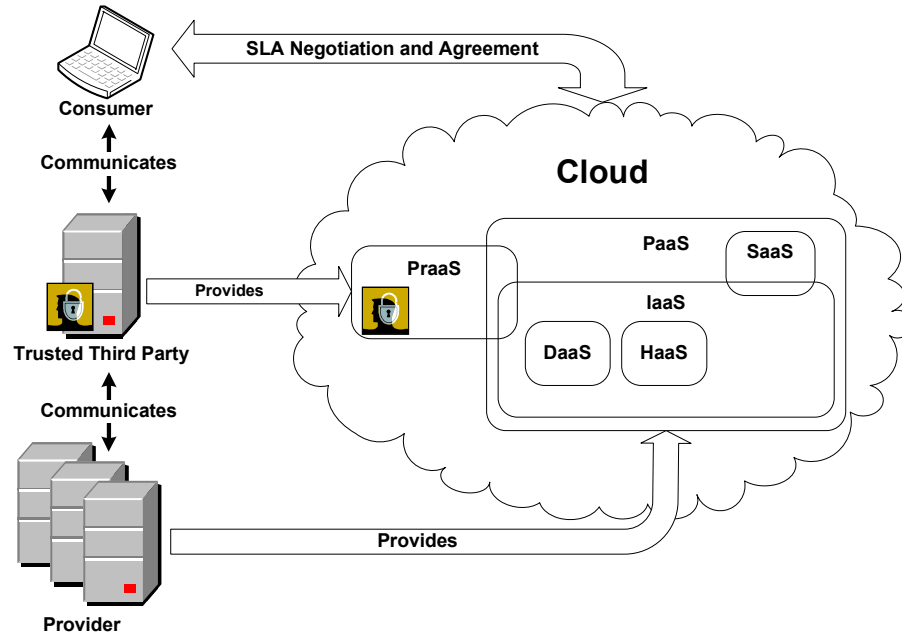


Fig. 3. Privacy as a Service Deployment

5.1 Monitoring Privacy with Privacy as a Service

We propose a new Privacy as a Service (PraaS) to handle the monitoring of the privacy agreement between the Cloud consumer and provider. The PraaS will ensure the agreement that was reached between Cloud consumer and provider is being adhered to. The PraaS must be created and introduced into the Cloud by a trusted third party (TTP). The PraaS will be monitoring the situation with regards to privacy within the Cloud, as well as making decisions that affect both the Cloud consumer and provider. Trust in the provider of the PraaS is required in order to ensure both the Cloud consumer and provider that this monitoring and decision making process is being done without bias. This TTP could be an established standard organization, such as the OECD [27] or World Wide Web Consortium (W3C) [28], a public sector organization, or a not-for-profit organization. The PraaS is illustrated in Figure 3.

Once the PraaS is introduced into the Cloud, it will be required to monitor the system. Monitoring in a Cloud environment is more difficult and complex than compared to an enterprise application due to the large amounts of data created over a distributed environment [29]. Compounding the problem is the lack of standardization in Cloud Computing [6]. The implementation of a monitoring system within PraaS is outside the scope of this work, but there are several existing Cloud monitoring tools which can provide examples of how such a system can be done [30] [31]. As many copies of the PraaS as required to efficiently monitor the system will be deployed in the Cloud. This will be done to avoid the PraaS becoming a single point of failure and to avoid creating a bottleneck in the Cloud.

5.2 Legislation and the Public Sector

If a privacy violation is detected by the PraaS, both the Cloud consumer and provider should be informed. The TTP that deploys the PraaS can perform this action, as shown in Figure 3. The communication can be handled automatically by the PraaS, or another type of service developed by the TTP. If no quick resolution to the problem can be found, or the violation is severe, enforcement will be required. Enforcement of any technology that operates over networks such as the Internet is a challenging problem. In order to provide accountability, governing bodies must create legislation addressing privacy in a Cloud environment. This legislation would assist the PraaS by providing tools for enforcement, such as an enforcement body to report to and punishments for infractions. Legislation is notoriously late when addressing problems in technology. Further research will be required to compare the differences between legislation that has already been adopted around the world, and what areas and topics still require new laws. Creating and abiding by legislation over a distributed environment is a massive topic, as different countries have different, and often conflicting laws. These problems can be compounded if information is sent across international borders. As such, determining exactly what legislation is required falls outside the scope of this work, and is an ongoing field of research [32].

6 Related Work

Research into privacy for Cloud Computing is still in its infancy, and as such, there are no set guidelines or benchmarks. HP Labs Singapore has recently begun work on a service called Trust Cloud [33] to address concerns of data protection and security. The goal of this Trust Cloud is to monitor any information or file a user places in the Cloud and notify that user if any of their data has been accessed, moved or modified by the Cloud provider [33]. The monitoring proposed by HP's Trust Cloud project is similar to the monitoring required in this work, however there are currently no details available on how HP achieves this monitoring as it is still in early research.

The idea of Privacy as a Service has been mentioned before in a few works, but never in the same context as presented here. Itani, Kayssi and Chehab [34] presented Privacy as a Service to provide data storage and processing in Cloud Computing architectures. This approach treats privacy as a strictly encryption based problem, where Privacy as a Service provides secure storage and processing of users' confidential data by utilizing tamper-proof cryptographic coprocessors. Private information is divided into three levels: full trust, compliance-based trust, and no trust. At the full trust level, there is no encryption applied to the data. At the compliance-based trust level, encryption is applied by the Cloud provider once the data has been sent by the consumer. At the no trust level, encryption is handled by a trusted third party before it is sent to the Cloud provider. This approach is very different from the approach presented in this work, as no privacy conditions are set, only varying methods of encryption are utilized.

There is another example of Privacy as a Service [35], however this work details a framework, service, model and algorithm to address shortcomings in social platforms, specifically Facebook. This work did not attempt to provide privacy in a greater Cloud Computing environment, and shares very little with our work.

The work by Pearson [36] discusses the importance of considering privacy while designing a Cloud Computing environment. This work discusses the privacy challenges presented by Cloud Computing environments, and outlines nine privacy principles which closely resemble the eight privacy principles used by the OECD. It provides guidelines to follow to mitigate the risks of privacy, but does not provide a framework for protecting privacy.

7 Conclusions & Future Work

In order for computer software and hardware systems to evolve successfully, new technologies must not only be created, but also grow and gain the acceptance of consumers. Cloud Computing represents a large step forward in the evolution of software and hardware. Consumers are no longer forced to install their own copy on a single machine. With Cloud Computing, a consumer will have access to the same application using the same data from virtually anywhere. There are several issues Cloud Computing must address before it can truly become a widely accepted technology. Privacy is one of the biggest unaddressed issues Cloud Computing currently faces. It is important for Cloud Computing to gain a wider acceptance in order to take advantage of the many economic and environmental benefits it provides. The public sector also has a role to play in Cloud Computing privacy, by providing the legislation required for enforcement of privacy, and by providing an example by adopting the Cloud itself.

In this work we introduced the first steps in creating a privacy solution for Cloud Computing. This solution involved identifying key aspects of privacy that must be represented in a formal agreement between the Cloud consumer and provider. These privacy aspects are derived from privacy legislation used throughout the world [24]. We created quality of service levels from these privacy constraints, in order to incorporate privacy into the Cloud SLA. This novel solution provides a mechanism for creating a contract between Cloud consumer and provider that outlines how private information can be used. To the best of our knowledge, this is the first attempt at incorporating privacy agreement terms into a Cloud Computing SLA.

Finally this work introduced a new Privacy as a Service (PaaS). This PaaS is hosted by a trusted third party and tasked with the job of both monitoring for privacy violations and creating accountability through enforcement. Enforcement is only effective when coupled with appropriate legislation, which also must be addressed.

This work is intended to be the first of many steps towards providing privacy protection in Cloud Computing environments. As such, there are many future directions for this research:

- The proposed SLA privacy extension will be further researched to determine if any refinements need to be done. This can be carried out through different case studies, with interest taken into how the current design of the SLA privacy terms handles different scenarios.
- A study into the different ways Cloud SLAs are formatted will be completed. If necessary, a standard SLA will be selected for the privacy extension.
- More studies will be conducted into current legislation from around the world that pertains to Cloud Computing and privacy. This research will help to

identify how governments are approaching the problem, and to better develop monitoring and enforcement of the PraaS.

- Further research will be done to identify the best monitoring solution for the PraaS. There are many Cloud monitoring options available, and the most efficient and effective will be selected.
- The PraaS will first be developed in a laboratory setting, where simulations can be run to test its effectiveness and performance. Following this, the ultimate plan for this research is to implement the PraaS in a real world scenario. There is the chance that the state of local legislation at the time of the PraaS development will not allow the PraaS to provide effective enforcement. In this case, the PraaS will focus on monitoring and reporting, allowing for the future expansion of enforcement when possible.

In order for Cloud Computing to achieve widespread long term success, fundamental issues such as privacy must be addressed. This work takes the first steps towards this goal, with the hope that it will lead towards greater success and acceptance of Cloud Computing technology.

References

1. Google Mail - Gmail: Email from Google, <http://mail.google.com>.
2. Google Docs - Online Documents, Spreadsheets, Presentations, Surveys, File Storage and More, <http://docs.google.com>.
3. SolarWinds Cloud Survey Press Release, http://www.solarwinds.com/Company/Newsroom/Press_Releases/Years/2010/21474837235.aspx.
4. Yee, G.: Privacy Protection for E-Services, IGI Publishing, Hershey, PA, USA (2006).
5. Ryan, M.: Cloud Computing Privacy Concerns on Our Doorstep. *Communications of the ACM*. 54(1), pp. 36--38 (2011).
6. Patel, P., Ranabahu, A., Sheth, A.: Service Level Agreement in Cloud Computing, http://knoesis.wright.edu/library/download/OOPSLA_Cloud_wsla_v3.pdf.
7. Gartner Estimates ICT Industry Accounts for 2 Percent of Global CO2 Emissions Press Release, <http://www.gartner.com/it/page.jsp?id=503867>.
8. Microsoft Corporation, <http://www.microsoft.com>.
9. Accenture, <http://www.accenture.com>.
10. WSP Environment & Energy, <http://www.wspenvironmental.com>.
11. Cloud Computing and Sustainability: The Environmental Benefits of Moving to the Cloud, http://www.microsoft.com/click/services/Redirect2.ashx?CR_EAC=300012377.
12. Allison, D., EL Yamany, H., Capretz, M.: A Privacy Service for Comparison of Privacy and Trust Policies within SOA. In: Gupta, M., Walp, J., Sharman, R. (eds.) *Strategic and Practical Approaches for Information Security Governance: Technologies and Applied Solutions*, IGI Global, New York, NY, USA (2011).
13. Zhou, M., Zhang, R., Xie, W., Qian, W., Zhou, A.: Security and Privacy in Cloud Computing: A Survey. In: *Proceedings of the Sixth International Conference on Semantics, Knowledge and Grids*. Ningbo, China (2010).
14. Sun, W., Zhang, K., Chen, S., Zhang, X., Liang, H.: Software as a Service: An Integration Perspective. In: Krämer, B., Lin, K., Narasimhan, P. (eds.) *Service-Oriented Computing - ICSSOC 2007*. LNCS, vol. 4749, pp. 558--569. Springer, Heidelberg (2007).
15. Durkee, D.: Why Cloud Computing Will Never Be Free. *Queue*. 8(4), (2010).

16. Wang, L., Tao, J., Kunze, M.: Scientific Cloud Computing: Early Definition and Experience. In: Proceedings of the 10th IEEE International Conference on High Performance Computing and Communication, pp. 825--830. Dalian, China (2008).
17. Truong, H., Dustdar, S.: On Analyzing and Specifying Concerns for Data as a Service. In: Proceedings of the 2009 IEEE Asia-Pacific Services Computing Conference, pp. 87-94. Biopolis, Singapore (2009).
18. Lawton, G.: Developing Software Online with Platform-as-a-Service Technology. *Computer*. 41(6), pp. 13--15 (2008).
19. Nelson, M.: The Cloud, the Crowd, and Public Policy. *Issues in Science and Technology*. 25(4). pp. 71--76 (2009).
20. Cooley, T.: A Treatise on the Law of Torts or the Wrongs Which Arise Independent of Contract, 2d ed. Callaghan & Co., Chicago, IL, USA (1888).
21. Kanneganti, R., Chodavarapu, P.: SOA Security. Manning Publications Co., Greenwich, CT, USA (2008).
22. Comuzzi, M., Kotsokalis, C., Spanoudakis, G., Yahyapour, R.: Establishing and Monitoring SLAs in Complex Service Based Systems. In: Proceedings of the 2009 IEEE International Conference on Web Services, pp. 783--790. Los Angeles, CA, USA (2009).
23. Organisation for Economic Co-operation and Development.: OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, http://www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html.
24. Cavoukian, A., Hamilton, T.: The Privacy Payoff: How Successful Businesses Build Customer Trust, McGraw-Hill Ryerson Limited, Whitby, ON, Canada (2002).
25. Allison, D., EL Yamany, H., Capretz, M.: Metamodel for Privacy Policies within SOA. In: The Proceedings of the 5th IEEE International Workshop on Software Engineering for Secure Systems in conjunction with the 31st IEEE International Conference of Software Engineering. Vancouver, BC, Canada (2009).
26. EL Yamany, H., Capretz, M., Allison, D.: Quality of Security Service for Web Services within SOA. The Proceedings of the 2009 IEEE International Conference on Cloud Computing, Los Angeles, CA, USA (2009).
27. Organisation for Economic Co-operation and Development, <http://www.oecd.org>.
28. World Wide Web Consortium, <http://www.w3.org>.
29. Mathew, J.: Monitoring Applications in the Cloud, http://www.infosysblogs.com/cloudcomputing/2010/05/cloud_based_monitoring_tools.html.
30. Nagios - The Industry Standard in IT Infrastructure Monitoring, <http://www.nagios.org>.
31. Hyperic - Systems Monitoring, Server Monitoring & Systems Management Software, <http://www.hyperic.com>.
32. Pearson, S., Charlesworth, A.: Accountability as a Way Forward for Privacy Protection in the Cloud. In: Proceedings of the 1st International Conference on Cloud Computing, pp. 131--144. Beijing, China (2009).
33. HP Labs Singapore Aims for the Clouds: Research Focuses on Customer Co-Innovation and Democratizing the Cloud, http://www.hpl.hp.com/news/2011/apr-jun/hp_labs_singapore.html.
34. Itani, W., Kayssi, A., Chehab, A.: Privacy as a Service: Privacy-Aware Data Storage and Processing in Cloud Computing Architectures. In: Proceedings of the Eight IEEE International Conference on Dependable, Autonomic and Secure Computing, pp. 711-716. Chengdu, China (2009).
35. Maximilien, E., Grandison, T., Sun, T., Richardson, D., Guo, S., Liu, K.: Privacy-as-a-Service: Models, Algorithms, and Results on the Facebook Platform. In: Proceedings of Web 2.0 Security and Privacy. Oakland, CA, USA (2009).
36. Pearson, S.: Talking Account of Privacy when Designing Cloud Computing Services. In: Proceedings of the Workshop on Software Engineering Challenges in Cloud Computing in conjunction with the 31st IEEE International Conference of Software Engineering. Vancouver, BC, Canada (2009).