

NRC Publications Archive Archives des publications du CNRC

Measuring Privacy Protection in Web Services

Yee, George

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version.
/ La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

Proceedings of the IEEE International Conference on Web Services 2006 (ICWS 2006), 2006

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=ce4cba76-5667-49e8-8a39-b69d46c085e>
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=ce4cba76-5667-49e8-8a39-b69d46c085eb>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.



National Research
Council Canada

Conseil national
de recherches Canada

Institute for
Information Technology

Institut de technologie
de l'information

NRC - CNRC

*Measuring Privacy Protection in Web Services **

Yee, G.
November 2006

* published in the Proceedings of the IEEE International Conference on Web Services 2006 (ICWS 2006). Chicago, Illinois, USA. September 18-22, 2006. NRC 48734.

Copyright 2006 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables from this report, provided that the source of such material is fully acknowledged.

Measuring Privacy Protection in Web Services¹

George Yee

Institute for Information Technology

National Research Council Canada

george.yee@nrc.ca

Abstract

The growth of the Internet has been accompanied by the growth of web services (e.g. e-commerce, e-health) leading to the need to protect the personal privacy of web service users. However, it is also important to be able to measure a web service in terms of how well it protects personal privacy. Such a capability would benefit both users and developers. Users would benefit from being able to choose (assuming that such measures were made public) the service that has the greatest ability to protect user privacy (this would in turn encourage web service providers to pay more attention to privacy). Developers would benefit by being able to incrementally measure and modify their services during development until certain target levels of privacy protection are reached. This paper presents an approach for measuring how well a web service protects personal privacy and illustrates the approach with an example.

1. Introduction

This work considers web services to be: a) web-based services that employ XML (eXtensible Markup Language), WSDL (Web Service Definition Language), SOAP (Simple Object Access Protocol), and UDDI (Universal Description, Discovery, and Integration) in a service oriented architecture (SOA) [1], and b) existing and previous generations of web-based applications that involve web browsers interacting with web servers that do not employ XML, WSDL, SOAP or UDDI. This work applies to all web services so described.

Numerous web services targeting consumers have accompanied the rapid growth of the Internet. Web services are available for banking, shopping, learning, healthcare, and Government Online. However, most of these services require a consumer's personal information in one form or another, leading to concerns

over privacy. For web services to be successful, privacy must be protected. Various approaches have been used to protect personal information, including data anonymization [2,3] and pseudonym technology [4]. Approaches for privacy protection that are in the research stage include treating privacy protection as an access problem and then bringing the tools of access control to bear for privacy control [5], treating privacy protection as a privacy rights management problem using the techniques of digital rights management [6], and considering privacy protection as a privacy policy compliance problem, verifying compliance with secure logs [7].

It is also important to measure how well a web service protects consumer privacy. Suppose such measures are calculated for similar web services A, B, and C and made available to consumers. This leads to the following benefits. If the consumer has to choose one service from among A, B, and C, then the measures can help the consumer decide which service to select (probably the service that has the highest level of privacy protection). In addition, the fact that consumers have access to these measures may encourage service providers to pay more attention to protecting consumer privacy and result in higher levels of consumer trust and acceptance of web services. Alternatively, web service developers can use measures of how well a service protects consumer privacy to develop services that meet pre-defined goals of privacy protection. Pre-defined levels of the measures could be expressed as development requirements. The measures could then be evaluated for incremental versions of a service until the pre-defined levels are achieved.

The objectives of this paper are to a) define measures of how well a web service protects consumer privacy, b) show how the measures are calculated, and c) illustrate the calculation and application of the measures using a web services example.

The rest of this paper is organized as follows. Section 2 defines the measures. Section 3 shows how to calculate the measures. Section 4 illustrates the

calculation and application of the measures. Section 5 discusses related work. Section 6 gives our conclusions and directions for future research.

2. Measures of privacy protection

2.1. Privacy

In order to define measures of how well a web service protects consumer privacy, it is necessary first to examine the nature of personal privacy. As defined by Goldberg et. al. in 1997 [8], privacy refers to the ability of individuals to *control* the collection, retention, and distribution of information about themselves. This leads to the following definitions for this work.

DEFINITION 1: *Privacy* refers to the ability of individuals to *control* the collection, use, retention, and distribution of information about themselves.

DEFINITION 2: A service’s *protection of user privacy* refers to the service’s use of provisions to give the user control over the service’s collection, retention, and distribution of information about the user.

DEFINITION 3: A *measure* of a service’s protection of user privacy is a numerical value that indicates the degree of the user’s control (or some aspect of that control) over the service’s collection, retention, and distribution of information about the user.

Definition 1 is the same as given by Goldberg et. al. except that it also includes “use”. To see that “use” is needed, consider, for example, that one may agree to give out one’s credit card number to pay for one’s book (from an online bookstore) but not to pay for someone else’s book. The “provisions” in Definition 2, refer to whatever means or technologies are needed to give the user the required control (uphold the user’s privacy), e.g. privacy preference languages, policy negotiation mechanisms, access control mechanisms, policy compliance mechanisms. These provisions depend on the nature of the control required by the user.

It follows from Definition 2, that if the service provider is to make provisions to uphold the user’s privacy, it needs to know how the user wishes to control her (“her” and “she” are used here to stand for both sexes) personal information. Thus, there must be a means of communicating the nature of this control, from the user to the service provider. This communication is normally carried out using a statement of privacy preferences called a *privacy*

policy. Figure 1 is an example of a consumer privacy policy for e-learning from [9]. In Figure 1, each item of information about the user corresponds to a “privacy rule” that spells out how the item is to be collected, used (purpose), retained, and distributed (disclose-to). Figure 2 illustrates the use of a privacy policy to express how the user wishes to control her private information. The service provider would have to agree to uphold the user’s privacy policy before it can receive any of the user’s private information. Where the service provider does not agree to the user’s policy, the user can negotiate with the provider [10, 11] until there is agreement, or the user can try a different provider.

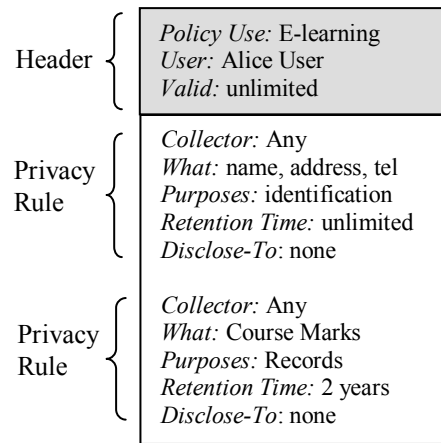


Figure 1. Example privacy policy for e-learning

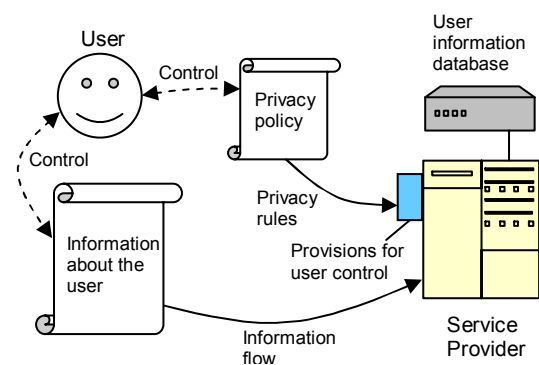


Figure 2. Using a privacy policy to express user control over private information

2.2. Privacy violations

Once the provider has agreed to uphold the user’s privacy policy and is in possession of the user’s private

data, the user's privacy is fully protected, provided there are no violations of the user's privacy policy. Such violations can be classified as internal and external violations:

- Internal Violations (IV): The provider or employees of the provider may be dishonest and violate the policy for their own gain (e.g. selling the user's contact information to commercial spammers). These may also be called insider attacks.
- External Violations (EV): The provider's systems that provide the service or store private data may be attacked by malicious parties outside the provider's organization compromising the user's private information (e.g. Trojan horse attack).

Thus, measuring how well a service provider protects privacy involves looking at what provisions are in place to prevent IV and EV. Let M denote the measure of how well a service provider protects consumer privacy. Measure M will contain two components: one component, m_i , to account for the provisions used against IV and the other component, m_e to account for the provisions used against EV. In other words, M is a matrix expressed as

$$M = (m_i, m_e).$$

In the case of provisions against IV, it is necessary to assume that the provider and its employees are dishonest unless proven otherwise, in order to identify the provisions (if they are honest there would be no need for such provisions). The following provisions aim to prevent IV or lessen the probability of it occurring:

- Use of a privacy policy compliance system (PPCS) (e.g. [7, 12]) that automatically ensures that the user's privacy policy is not violated,
- Use of a cryptographically secure log (this log can be later inspected to check for policy violations) to record each provider action involving the user's private data,
- Use of employee background checks when they are hired to try to exclude dishonest people from the provider's organization,
- Use of reputation mechanisms to record and indicate the past performance of the provider in terms of integrity (e.g. Better Business Bureau),
- Use of seals of approval that attest to the fact that the provider has undergone and passed rigorous inspections of its business processes (e.g. ISO 9001: 2000 [13]).

This list is of course not exhaustive. A service provider may employ none, one, or more than one of these or

other provisions. In order to avoid possible ineffective use of these provisions, it is recommended that a standards body such as the Organization for the Advancement of Structured Information Standards (OASIS, <http://www.oasis-open.org/home/index.php>) or the International Organization for Standardization (ISO, <http://www.iso.org/iso/en/ISOOnline.frontpage>) study and recommend combinations of these provisions as being effective against IV. In addition, the standards body would provide a percentage rating of the effectiveness of each combination. Let p_j denote the percentage effectiveness of combination j . Then for a service provider that has implemented combination k ,

$$m_i = p_k, \quad 0 \leq p_k \leq 1$$

In the case of provisions against EV, the question to ask is: "What are possible external violations against a user's privacy?" Whereas internal violations have to do with not upholding the user's privacy policy, external violations are concerned with gaps in security that allow an attacker to access the user's private information. Hence it is recommended that a special security threat analysis [14], oriented towards discovering security weaknesses that would jeopardize the user's private information, be carried out. Suppose that such an analysis identified n such security weaknesses but that effective security provisions (or countermeasures) are in place for q of the weaknesses. Then for a service provider with such analysis results,

$$m_e = q/n, \quad \text{if } n > 0, \text{ so that } 0 \leq m_e \leq 1 \\ = 1, \quad \text{if } n = 0.$$

Substituting the values for m_i and m_e into the equation for M ,

$$M = (p_k, q/n), \quad \text{if } n > 0 \\ = (p_k, 1), \quad \text{if } n = 0.$$

In practice, m_i and m_e may be more visible to consumers expressed on a scale of 1 to 10. Therefore, rather than using M directly, it is recommended that M_{10} be used to measure how well a service provider protects privacy, where

$$M_{10} = (10.p_k, 10.q/n), \quad \text{if } n > 0 \\ = (10.p_k, 10), \quad \text{if } n = 0$$

By setting minimum acceptable thresholds t_i and t_e for $10.m_i$ and $10.m_e$ respectively (thresholds above which the corresponding provisions for IV and EV are sufficient for privacy protection), a shaded region is defined in Figure 3 in which the service provider "passes" M_{10} , i.e. where the privacy protection capability of the provider is acceptable.

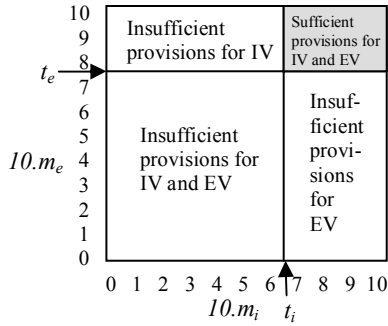


Figure 3. Region (shaded) in which a service provider's provisions for IV and EV are acceptable

2.3. Multi-provider services

The provider of a web service may make use of other providers in offering its service. For example, an online book store (e.g. Amazon.com) may make use of an online payment service (e.g. Paypal.com) and a shipping service (e.g. fedex.com) to sell a book. For the sake of exposition, the *primary* service is the service that the consumer chooses. *Secondary* services are services that the primary service makes use of in order to complete itself and are provided by other providers. Such a primary service is called a *multi-provider* service. In order to evaluate such a primary service for its privacy protection effectiveness, it is necessary to evaluate in addition all secondary services of the primary service that receive the user's private information. Such secondary services are identified in the user's privacy policy under "disclosed-to". Of course, all such secondary services have to agree to uphold the user's privacy policy. This gives rise to the following definition.

DEFINITION 4: A *multi-provider service* passes M_{10} evaluation if and only if the p_k and q/n components of M_{10} are each above or equal to their respective thresholds for the primary service and each of its secondary services that receives the user's private information.

This definition reflects consumer wishes since the consumer would not only want the primary service to respect her privacy preferences but would also want each secondary service of the primary service to respect them as well.

3. Calculation of the measures

3.1. Calculation of m_i

As mentioned in Section 2, a standards body is given the job of standardizing and recommending various combinations of provisions for combating IV. In addition, the standards body is to provide a percentage rating of the effectiveness of each recommended combination. These are not difficult tasks for subject matter experts that are part of every standards body. Table 1 gives some examples of such combinations along with their ratings.

Table 1. Example IV provision combinations

Comb. number	Description	Effective -ness Rating (p_k)
1	PPCS only	95%
2	Secure log only	60%
3	Secure log, employee screening, reputation mechanism	70%
4	Secure log, employee screening, reputation mechanism, seals of approval	80%

Note that in Table 1, the use of a PPCS has such a high effectiveness rating that if used, no other provision is really needed. In addition, the lower ratings may be usable for certain services where the demand for privacy is not very high, e.g. information services that only require a name.

3.2. Calculation of m_e

The calculation of m_e requires a threat analysis of security weaknesses in the service provider's systems that could allow an outside attacker to have access to the user's private information. An overview of threat analysis follows.

Threat analysis or threat modeling is a method for systematically assessing and documenting the security risks associated with a system [14]. The results can help development teams identify the strengths and weaknesses of the system and serve as a basis for investigations into vulnerabilities and required mitigation. Threat modeling involves understanding the adversary's goals in attacking the system based on the system's assets of interest. It is predicated on that fact that an adversary cannot attack a system without a way of supplying it with data or otherwise accessing it. In addition, an adversary will only attack a system if it has some assets of interest. The following threat modeling terminology is selected from [14]:

- *Attack path:* A sequence of conditions in a threat tree that must be met for an attack goal

(threat) to be achieved. A valid attack path (one with no mitigated conditions) is a vulnerability.

- *Threat*: The adversary's goals, or what an adversary might try to do to a system. Threats to a system always exist, regardless of mitigation.
- *Threat Tree or Attack Tree*: An analysis tool that describes the attack paths for a particular threat. A threat tree is comprised of hierarchical conditions and allows the threat's mitigation (or lack thereof) to be characterized. The root of the threat tree is the threat to which the tree corresponds.

This work adapts the method for threat modeling given in [14] for use in calculating m_e . The steps from [14] are:

1. Create attack trees for the system.
2. Apply weights to the leaves.
3. Prune the tree so that only exploitable leaves remain.
4. Generate corresponding countermeasures.
5. Optimize countermeasure options.

However, the above steps are oriented towards developing secure systems. For calculating m_e , the above steps are modified to (descriptions follow):

1. Identify threats on the user's data.
2. Create attack trees for the system.
3. Apply weights to the leaves.
4. Prune the tree so that only exploitable leaves remain. Count the number of such leaves or vulnerabilities (this gives the n for m_e).
5. Determine if countermeasures are in place for the vulnerabilities found in step 4. Count the number of these vulnerabilities so mitigated (this gives the q for m_e).

Step 1: Identify threats on the user's data.

In this step, examine the architecture and all available details of the system and enumerate possible threats on the user's data. Represent the system pictorially to get the big picture. Disregard any existing security countermeasures – they will be accounted for in step 5. This step requires experience and imagination and may involve confirming details with the system's developers.

Step 2: Create attack trees for the system.

Corresponding to each threat identified in step 1, systematically create an attack tree, by putting yourself in the adversary's place in finding the weak points in the system and the paths which will lead to realizing the threat. This analysis terminates in a series of vulnerability leaves for each attack tree. (In this work, each attack tree is represented by hierarchical indented

headings rather than pictorially, which takes up too much space and becomes unwieldy).

Step 3: Apply weights to the leaves.

For each leaf, assign qualitative values (high, medium, low) for adversary risk, impediment to access, and cost. For example, an adversary sending an email containing a virus attachment has low risk (probability of being identified is low), medium access (probability of the victim not opening the attachment and unleashing the virus is medium), and low cost (cost to the adversary to create the virus email is low).

Step 4: Prune the tree so that only exploitable leaves remain. Count the number of such leaves or vulnerabilities.

Countermeasures are required only for attacks that meet an adversary's objectives. Such attacks must match the adversary's capabilities and offer an adequate return. Prune the tree of leaves that fail these criteria. For example, if a particular attack requires 2^{256} bytes of computer memory, it could safely be pruned as beyond the resources of any adversary [14]. As another example, if an attack requires access to a heavily guarded military installation, the risk may be too great for most adversaries. After pruning the tree, count the number n of exploitable leaves or vulnerabilities that remain.

Step 5: Determine if countermeasures are in place for the vulnerabilities found in step 4. Count the number of these vulnerabilities so mitigated.

Examine what countermeasures are in place for the vulnerabilities found in step 4 and count the number q of vulnerabilities mitigated by the countermeasures. This requires experience and knowledge of security countermeasures.

After performing the above steps, both q and n are available for calculating m_e .

4. Application example

Consider a web service, Easy123Drugs.com, that is an online drug store (e.g. Walgreens.com). Easy123Drugs is a multi-provider service that makes use of two business web services: an online payment service PayAsYouLikeIt.com (e.g. Paypal.com) and an accounting service AccountingAsNeeded.com (e.g. cbiz.com). Suppose Easy123Drugs, PayAsYouLikeIt, and AccountingAsNeeded (all fictitious names with no hits on Google) are all web services that are based on the Service Oriented Architecture [1], employing XML-based protocols (not necessarily the case for the

real life examples cited here). Due to space limitations in this paper, the details regarding UDDI lookup and service binding via SOAP and WSDL [1] will not be described here. It is assumed that these initialization steps occur as required. Figure 4 shows the network architecture of these services after service lookup and binding have occurred. The dashed lines in Figure 4 indicate logical communication channels.

Table 2 shows the consumer's private information required by each service along with their provisions against IV (per recommendations from a standards body based on their willingness to pay for such provisions). The values of m_i as rated by the standards body are also shown (from Table 1). The consumer provides her private information to Easy123Drugs once her privacy policy has been accepted and agreed to by all the services. Easy123Drugs then discloses this information to PayAsYouLikeIt and Accounting-AsNeeded according to the consumer's privacy policy.

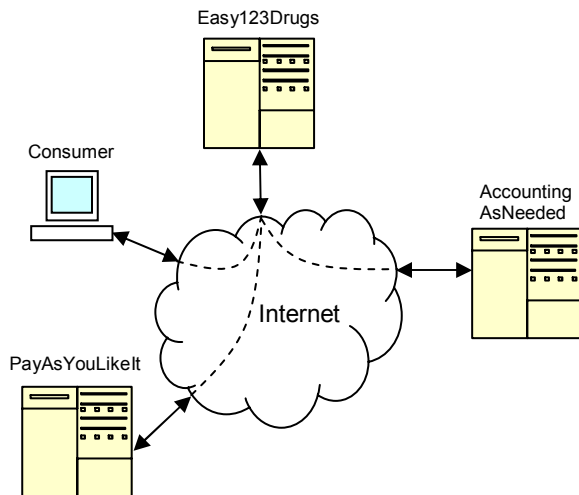


Figure 4. Network architecture of Easy123Drugs service

Table 2. Private information required and provisions against IV

Web Service	Private Information Required	Provisions Against IV	m_i
Easy123Drugs	Consumer's name, drug name, Doctor's name and authorization, consumer's address	PPCS	.95

PayAsYouLikeIt	Consumer's name, credit card details	Secure log	.60
AccountingAsNeeded	Consumer's name, drug name, Doctor's name, quantity of drug sold, price paid by consumer, consumer's address	Secure log, employee screening, reputation mechanism, seals of approval	.80

The threat analysis described in section 3.2 is now applied to calculate m_e for each service. Assume that each service stores the user's private data in a database and that the external threats to the user's private data are the same for each service (not an unrealistic assumption). It is then possible to do one threat analysis that applies to all three services. However, the countermeasures in place are likely to be different for each service. Following the steps in section 3.2,

Step 1: Using Figure 4 to visualize possible threats against the user's data, the threat that is at the root of all possible threats is: "external attacker compromises the user's data".

Steps 2 and 3: The attack tree and weights are as follows. The weights are expressed as a triple (risk, access, cost) where the values for risk (to the attacker), access (impediments to access for the attacker), and cost (cost to the attacker) can be high (H), medium (M), or low (L).

1. External attacker compromises the user's data.
 - 1.1. Attacker steals the user's data.
 - 1.1.1. Attacker launches a man-in-the-middle attack on a communication channel to eavesdrop. (L, L, L)
 - 1.1.2. Attacker launches a Trojan horse attack on a provider's system. (L, L, L)
 - 1.1.3. Attacker launches a phishing attack on the user. (L, L, M)
 - 1.1.4. Attacker uses social engineering to deceive a provider staff member into giving out the user's data. (M, M, L)
 - 1.1.5. Attacker breaks into a provider's premises to steal the user's data. (H, H, M)
 - 1.1.6. Attacker mugs a provider employee and steals her access card to enter a provider's

premises and steal the user's data. (H, H, L)

- 1.2. Attacker modifies the user's data.
 - 1.2.1. Attacker launches a man-in-the-middle attack on a communication channel to modify the user's data. (L, L, L)
 - 1.2.2. Attacker launches a virus attack on a provider's system. (L, L, L)
 - 1.2.3. Attacker uses social engineering to deceive a provider staff member into giving the attacker access to modify the user's data. (M, M, L)
 - 1.2.4. Attacker breaks into a provider's premises to modify the user's data. (H, H, M)
 - 1.2.5. Attacker mugs a provider employee and steals her access card to enter a provider's premises and modify the user's data. (H, H, L)

Step 4: The attack tree can be pruned by removing attack paths that are weighted with at least two H's. Applying this criterion, removes the attack paths (1, 1.1, 1.1.5), (1, 1.1, 1.1.6), (1, 1.2, 1.2.4), and (1, 1.2, 1.2.5). This leaves 7 vulnerabilities which can be assigned to each provider as follows: Easy123Drugs gets the full $n=7$ vulnerabilities, PayAsYouLikeIt gets $n=6$ vulnerabilities since the phishing attack really only applies to Easy123Drugs, and AccountingAsNeeded gets $n=6$ vulnerabilities, again because the phishing attack doesn't apply to it. Note that the man-in-the-middle attack on a channel is double counted when it is considered a vulnerability for the provider at each end of the channel. However, this double counting is remedied by the countermeasure, which removes the vulnerability from both providers.

Step 5: Suppose that Easy123Drugs has countermeasures in place against all vulnerabilities except phishing (the exact nature of the countermeasures is not important for this example). Suppose also that PayAsYouLikeIt and AccountingAsNeeded have countermeasures in place against all vulnerabilities except social engineering. Therefore, $q = 6$ for Easy123Drugs, $q = 5$ for PayAsYouLikeIt, and $q = 5$ for AccountingAsNeeded.

The values of m_i , m_e , and M_{10} can now be obtained for each provider as given in Table 3. Plotting these results for minimum acceptable thresholds $t_i=8$ and $t_e=8$ according to Figure 3 gives Figure 5, which shows that all services pass M_{10} on an individual basis, i.e. individually has sufficient provisions for IV and EV, except for PayAsYouLikeIt.

Table 3. Results of privacy protection evaluation

Service	n	q	m_i	$m_e = q/n$	M_{10}
Easy123Drugs	7	6	.95	.86	(9.5, 8.6)
PayAsYouLikeIt	6	5	.60	.83	(6.0, 8.3)
AccountingAsNeeded	6	5	.80	.83	(8.0, 8.3)




Figure 5. Plots of M_{10} for the example services

This latter service is deficient in provisions for IV. For Easy123Drugs to pass M_{10} as a multi-provider service (from Definition 4), PayAsYouLikeIt would have to beef up its provisions for IV. Had threshold t_e been set to 9.0, no service would pass, individually or otherwise. In this case, development would need to install a countermeasure against phishing and employees would need to be trained to resist social engineering in the services that lacked these countermeasures. This shows that measures of privacy protection effectiveness can be used as a tool by development or management (for social engineering training) to achieve predefined goals of privacy protection.

5. Related work

The literature is very scarce on works dealing with measuring privacy protection or even just measuring privacy. In fact only one paper by Cook [15] that is very remotely related was found. In this paper, Cook devises a measure of memory protection based on a model of a protection system. The relationship of Cook's work to this work is that protecting memory can be applied to protecting privacy since private data is at some point in memory. However, other than this, Cook's work has little to do with this work. For example, Cook does not talk about privacy in a service's context nor does he consider IV. Of course, Cook's measure is totally different, designed for application to memory protection. Nevertheless, Cook's conclusion, that the measure "provides a useful

tool for the designers of operating systems and other software” and resulted in changes made to improve the system, supports the claim of this work that developers can use measures of privacy protection to meet privacy protection goals.

6. Conclusions and future research

This work has defined measures for assessing how well a web service protects a user’s private information and illustrated the application of the measures using a web services example. The measures serve at least two important functions: 1) they help the consumer to choose services that are more effective at protecting privacy, and 2) they let web service developers or managers know if more countermeasures are needed to achieve a predefined level of privacy protection effectiveness.

The privacy protection effectiveness of all web services can be calculated by a privacy protection authority (to ensure fairness) and made available to consumers. This could encourage providers to achieve higher levels of privacy protection and lead to greater consumer trust in web services. Having standards bodies rate and recommend provisions against IV may be a little controversial but it would be a good use of their resources.

Future research includes improving the procedure for threat analysis by automating it and making it more foolproof as well as investigating other possible measures of privacy protection effectiveness.

References

- [1] M. O’Neill et al, *Web Services Security*, McGraw-Hill/Osborne, 2003.
- [2] V.S. Iyengar, “Transforming Data to Satisfy Privacy Constraints”, *Proceedings, SIGKDD’02*, Edmonton, Alberta, 2002.
- [3] A. Kobsa, J. Schreck, “Privacy Through Pseudonymity in User-Adaptive Systems”, *ACM Transactions in Internet Technology*, Vol. 3, No. 2, pp. 149-183, May 2003.
- [4] R. Song, L. Korba, and G. Yee, “Pseudonym Technology for E-Services”, chapter in *Privacy Protection for E-Services*, edited by G. Yee, Idea Group, Inc., 2006.
- [5] C. Adams and K. Barbieri, “Privacy Enforcement in E-Services Environments”, chapter in *Privacy Protection for E-Services*, edited by G. Yee, Idea Group, Inc., 2006.
- [6] S. Kenny and L. Korba, “Adapting Digital Rights Management to Privacy Rights Management”, *Computers & Security*, Vol. 21, No. 7, November 2002, 648-664.
- [7] G. Yee, L. Korba, “Privacy Policy Compliance for Web Services”, *Proceedings, 2004 IEEE International Conference on Web Services (ICWS 2004)*, San Diego, California, July 6-9, 2004.
- [8] I. Goldberg, D. Wagner, and E. Brewer, “Privacy-Enhancing Technologies for the Internet”, *IEEE COMPCON’97*, pp. 103-109, 1997.
- [9] G. Yee, L. Korba, “Semi-Automatic Derivation and Use of Personal Privacy Policies in E-Business”, *International Journal of E-Business Research*, Vol. 1, No. 1, pp. 54-69, Idea Group Publishing, 2005.
- [10] G. Yee, L. Korba, “The Negotiation of Privacy Policies in Distance Education”, *Proceedings, 14th IRMA International Conference*, Philadelphia, Pennsylvania, May 18-21, 2003.
- [11] G. Yee, L. Korba, “Bilateral E-services Negotiation Under Uncertainty”, *Proceedings, The 2003 International Symposium on Applications and the Internet (SAINT2003)*, Orlando, Florida, Jan. 27-31, 2003.
- [12] F. Lategan, M. Olivier, “PrivGuard: A Model to Protect Private Information Based On Its Usage”, available as of Dec. 14, 2005 from: <http://mo.co.za/open/privgrd.pdf>
- [13] International Organization for Standardization, “Selection and use of the ISO 9000:2000 family of standards”, retrieved January 28, 2006 from: http://www.iso.org/iso/en/iso9000-14000/understand/selection_use/selection_use.html
- [14] C. Salter, O. Sami Saydjari, B. Schneier, J. Wallner, “Towards a Secure System Engineering Methodology”, *Proceedings of New Security Paradigms Workshop*, Sept. 1998.
- [15] D. Cook, “Measuring Memory Protection”, *Proceedings of the 3rd international conference on Software Engineering*, IEEE Press, May 1978.

¹ NRC Paper Number: NRC 48734