Privacy by Design

An Overview of
Privacy Enhancing Technologies

26th November 2008

Enterprise Privacy Group
Old Bank House
59, High Street
Odiham
Hants RG29 1LF
T: +44 (0)1256 702325
www.privacygroup.org

AN OVERVIEW OF PRIVACY ENHANCING TECHNOLOGIES

Introduction

The Enterprise Privacy Group (EPG) has developed a report on 'Privacy by Design' on behalf of the Information Commissioner's Office (ICO). The report, which was published by the ICO, explores how to apply privacy principles in organisations, and in particular how to promote the use of Privacy Enhancing Technologies (PETs).

EPG's work included more detailed research into PETs than could be reflected in the final paper, and in particular an overview of PETs prepared by Rae Harbird, a researcher at University College London's Department of Computer Science. In the interest of providing a more detailed description of PETs and the market for privacy technologies, this document is a full copy of that paper.

The views expressed in this paper do not necessarily reflect those of the Enterprise Privacy Group or its Member organisations, or the Information Commissioner's Office.

Defining Privacy Enhancing Technologies

There is no widely accepted definition for the term Privacy Enhancing Technologies (PETs) although most encapsulate similar principles; a PET is something that:

1. reduces or eliminates the risk of contravening privacy principles and legislation.
2. minimises the amount of data held about individuals.
3. empowers individuals to retain control of information about themselves at all times.

To illustrate this, the UK Information Commissioner's Office defines PETs[1] as:

> *"... any technology that exists to protect or enhance an individual's privacy, including facilitating individuals' access to their rights under the Data Protection Act 1998".*

The definition given by the European Commission[2] is similar but also includes the concept of using PETs at the design stage of new systems:

> *"The use of PETs can help to design information and communication systems and services in a way that minimises the collection and use of personal data and facilitates compliance with data protection rules. The use of PETs should result in making breaches of certain data protection rules more difficult and / or helping to detect them."*

History of PETs

The term 'privacy enhancing technology' was coined in 1995 when it appeared as the title of a ground breaking report commissioned by the Information and Privacy Commissioner of Ontario, Canada and the Dutch Data Protection Authority[3] [4]. Software which can be categorised as a PET pre-dates this definition by well over a decade. The first PET is acknowledged to be 'Mix networks' devised by David Chaum as means of achieving anonymous and unobservable communications over a network[5]. Indeed, his research forms the basis for some of the anonymous communication and email systems still in use.

[1] Information Commissioner's Office, UK. Data protection guidance note: Privacy enhancing technologies. http://tinyurl.com/56th6c
[2] European Union. Press release: Privacy enhancing technologies (PETs). http://tinyurl.com/6hcnrm
[3] Information and Privacy Commissioner, Ontario, Canada and Registratiekamer, The Netherlands. Privacy-enhancing technologies: the path to anonymity, volume 1. Technical report, 1995. http://tinyurl.com/5t6qsd
[4] Information and Privacy Commissioner, Ontario, Canada and Registratiekamer, The Netherlands. Privacy-enhancing technologies: the path to anonymity, volume 2. Technical report, 1995. http://tinyurl.com/59w67g
[5] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. Commun. ACM, 24(2):84–90, 1981. ISSN 0001-0782. doi: http://doi.acm.org/10.1145/358549.358563

Today, a more holistic view of PETs has evolved and there is a general understanding that PETs are congruent with good design objectives for any system or technology and can offer demonstrable business benefits and competitive advantages for organisations that adopt them. PETs should not be "bolted-on" to systems or technologies that are privacy-invasive. We must use privacy engineering so that privacy-related objectives are considered alongside business goals and privacy considerations addressed at every stage of the systems development life-cycle[6].

Finally, in recognition of the fact that information systems increasingly transcend borders and jurisdictions, there is a move towards the specification of a consistent, global privacy management standard that can be used to evaluate whether the privacy-related objectives of any system have been achieve. There are several ongoing initiatives in support of this activity through standards bodies such as the International Standards Organisation (ISO)[7], The International Security, Trust and Privacy Alliance, (ISTPA)[8], and the European Committee for Standardisation and Information Society Standardisation System (CEN/ISSS)[9].

Classification

In the same way as there is no widely accepted definition for the term PETs neither is there a recognised means of classification. Recently though, some studies have categorised PETs according to their main function as either privacy management or privacy protection tools[10] [11]. Other PET classifications[12] are much more comprehensive and detailed. Such classifications are useful and necessary but for the purposes of describing what PETs do and listing some of the products available simplification is more appropriate.

Privacy management tools

Privacy management tools, also known as 'transparency tools' enable the user to look at the procedures and practices used by those who are handling personally-identifiable information (PII). They may also advise users of the consequences of the information processing performed leading to an improved understanding of privacy-related issues. There are a limited number of tools in existence today that cater for either the enterprise or the end-user market. The World Wide Web Consortium, known as the W3C, produced P3P[13], a protocol designed to give browsing users more control of their personal information. It allows web servers to declare their privacy policies with respect to the information collected, enabling users to negotiate the release of their details. The usefulness of P3P tools has been criticised because there is no enforcement mechanism implying that the actions of a service provider may differ from the intention stated in their P3P policy[14]. Further development of P3P has been suspended for now because none of the major browser vendors have integrated it into their products although it is still an active source of investigation in the research community.

---

[6] Dr Steve Marsh, Dr Ian Brown, and Fayaz Khaki. Privacy engineering whitepaper. http://tinyurl.com/5zv9b3

[7] International Organisations for Standardisation (ISO) and the International Electro-Technical Commission (IEC). Work of the JTC-1. http://tinyurl.com/2w6jqp

[8] International Security Trust and Privacy Alliance (ISTPA). Privacy framework. http://tinyurl.com/6hyanz

[9] European Committee for Standardisation (CEN). Information Society Standardization System (ISSS). http://tinyurl.com/6x3cz4

[10] Lothar Fritsch. State of the art of privacy-enhancing technology (PET). Technical Report 1013. http://publ.nr.no/4589.

[11] The META Group. Privacy enhancing technologies. http://tinyurl.com/6h3qru

[12] Carlisle Adams. A classification for privacy techniques. *University of Ottawa Law and Technology Journal*, 3(1):35–52, 2006. http://tinyurl.com/67tguj

[13] World Wide Web Consortium (W3C). Platform for privacy preferences (P3P). URL http://www.w3.org/P3P/

[14] Serge Egelman, Lorrie Faith Cranor, and Abdur Chowdhury. An analysis of P3P-enabled web sites among top-20 search results. In *ICEC '06: Proceedings of the 8th international conference on Electronic commerce*, pages 197–207, New York, NY, USA, 2006. ACM. ISBN 1-59593-392-1. http://tinyurl.com/6sb8h5

On an enterprise-level, products usually aim to enforce legal compliance. IBM's Secure Perspective software[15] allows organisations to create and manage enforceable security policies using natural language. Hewlett Packard's Openview Select Identity enables corporations to manage users and their entitlements. This is particularly interesting because recent research at HP has shown that Select Identity could be extended with the capability to enforce privacy-aware data life-cycle management from collection to disposal[16].

Privacy protection tools

Privacy protection tools, also known as opacity tools, aim to hide the user's identity, minimise the personal data revealed and camouflage network connections such that the originating IP address is not revealed. By learning your IP address an observer may be able to pinpoint your geographic location to the nearest town or city or even uniquely identify your computer. Privacy protection tools may also authenticate transactions such as payments while making it impossible to trace a connection back to the user. Some of the software that falls into this category is described here; it does not represent a definitive list.

- **Anonymising tools:** Software in this category hides the IP address of the originator and, in the case of anonymous or pseudonymous mail, the source email address. Some 'anonymous remailers', such as Mixminion[17], employ sophisticated techniques which enable receivers to reply to messages. More generally, Tor[18] is a network of virtual tunnels on the Internet that individuals and groups can use to keep websites from tracking them, to connect to news sites, instant messaging services or similar network services when these are blocked by their Internet service providers or may be sensitive in nature. A Firefox add-on, the Torbutton, provides a way to securely and easily enable or disable the browser's use of Tor at the click of a mouse. A feature known as 'hidden services' lets users publish web sites and other services without needing to reveal the location of the site. Journalists use Tor to communicate more safely with whistleblowers and dissidents. Non-governmental organizations (NGOs) use Tor to allow their workers to connect to their home website while they are in a foreign country, without notifying everybody nearby that they are working with that organisation.
- **Anonymous or pseudonymous payment:** The concept behind anonymous payment is straightforward and usually works in this way: the user purchases a pre-paid card which is identified by a unique number. When the user makes a purchase at an online store, payment is retrieved from the anonymous cash provider using the unique number on the card. Successful examples of commercial pre-paid cards include paysafecard[19] in Europe.

[15] IBM. IBM Secure Perspective. http://tinyurl.com/6528k4
[16] Hewlett Packard. Privacy-aware identity lifecycle management. http://tinyurl.com/6cv4sq
[17] George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: A type III anonymous remailer. URL http://mixminion.net/
[18] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: anonymity online. URL http://www.torproject.org/index.html.en
[19] PaySafeCard. URL http://www.paysafecard.com

- **Information security tools:** There are a number of applications for the end-user that are sometimes categorised as PETs but are, in fact, information security tools. These tools are important to data protection and privacy but their primary goal is to prevent unauthorised access to systems, files or communications over a network. The difference is best illustrated in an example by Lorrie Faith Cranor[20]: I can use encryption within my browser to communicate with an e-commerce site and this will prevent someone eavesdropping on the network and retrieving my credit card details. It will not prevent the online store from collecting excessive information about me or misusing my details in other ways. Some of the tools in this category have been widely adopted because they have been built into browsers and other standard computer software. In particular, most web-server and browser software can encrypt communications using the TLS or SSL protocol and this feature has been a significant factor in increasing confidence in online banking and ecommerce services. Other applications that fall into this category include: firewalls, virus checkers and spam filters.

Overview of recent and ongoing research

As might be expected, research in the PETs field is wide ranging and this is reflected in the diverse subject matter presented at conferences like the annual Privacy Enhancing Technologies Symposium[21]. A selection of recent PET-related research is presented here covering subjects as diverse as privacy and identity management, attacks on privacy using large desensitised data sets, secure voting systems and some of the techniques that are being applied to potentially privacy-invasive technologies to make them less intrusive.

The amount of our personal information stored and processed by both commercial and public organisations gives rise to concern. In business terms, customer-related data is extremely valuable and it is becoming easier than ever to analyse for financial benefit. Individuals have very little choice in handing over personal information which is often excessive for the intended purpose when interacting with organisations online. Research into user-centric identity management (U-Idm) frameworks may represent a viable solution to this problem. In most U-Idm frameworks users manage their own personal information which is stored on a personal computer or handheld device that they control. The user regulates the release of their personal information to organisations as required. U-Idm could facilitate update of, say, address information to multiple parties or provide proof of age or proof of entitlement online without revealing unnecessary identifying details.

Microsoft has recently acquired Credentica's U-Prove technology[22] which exploits special cryptographic techniques enabling users to enforce data minimisation or prove certain characteristics. It is easy to detect misuse such as forgery and, in this circumstance, a transgressor's identity can be revealed. Microsoft intends to embed these features in its U-Idm software, Windows CardSpace.

[20] Lorrie Faith Cranor. The role of privacy enhancing technologies. In Considering Consumer Privacy: A Resource for Policymakers and Practitioners. Center for Democracy and Technology, edited by Paula J. Bruening, March 2003

[21] Nikita Borisov and Ian Goldberg, editors. *Privacy Enhancing Technologies, 8th International Symposium, PETS 2008, Leuven, Belgium, July 23-25, 2008, Proceedings*, volume 5134 of *Lecture Notes in Computer Science*, 2008. Springer. ISBN 978-3-540-70629-8

[22] Credentica. U-Prove technology. URL http://www.credentica.com/

There are a number of recent and current research initiatives in this area working on the technical, social, legal and usability aspects of the technologies. The European Commission is currently funding PrimeLife[23], a collaboration between industry and academia. One of its goals is to produce identity management solutions that are widely-used and available as open source products. Two other European projects are conducting research in related fields. PICOS[24] is investigating solutions from the perspective of someone needing to use U-Idm features when out in the community, perhaps using their mobile phone or other handheld device. EnCoRe[25], Ensuring Consent and Revocation, is a UK project examining solutions in the area of consent and revocation with respect to personal information. This project will, over the next three years, help businesses and Government adopt scalable, cost–effective and robust consent and revocation methods for controlling the use, storing, locating and sharing of personal data. These new mechanisms should permit users to retain fine-grained control of their personal details. It would mean being able to withdraw your personal information from a company when you cease to be a customer.

The secure release, management and control of personal information in cyberspace represents a huge challenge especially when we consider that these activities may well be initiated and mediated without human intervention. Adoption of privacy enhancing technologies to support such activities will depend upon the existence of standard ways to describe our personal data and the manner in which it may be used. To illustrate this imagine automatically signing up for an online library service: first, the library's web server requests specific items of personal information from the prospective customer's browser. The library's request also contains the promises, expressed as privacy policies, which the library makes with respect to the treatment of that personal information. The privacy policy must be flexible enough to capture concepts such as purpose of use and requests for consent. To achieve this first step alone the personal data fields and policies must be described in a way that is universally understood and this is generally achieved using tags called metadata.

Conversely, the device acting on behalf of the customer will evaluate the web server's request against the user's personal privacy preferences and reply with the information if appropriate. The user's personal data may also be accompanied by a set of conditions, known as obligations, covering such things as the length of time that the library may keep the information before deleting it or whether the user's consent is given for passing the information to third parties such as other local government departments. Work is underway in the research community, including in the PrimeLife and EnCoRE projects, to investigate the detailed requirements for policy languages to support the type of functionality described in this scenario and to encourage use and dissemination of standards through bodies such as W3C's Policy Languages INterest Group (PLING).

[23] PrimeLife. PrimeLife - bringing sustainable privacy and identity management to future networks and services. http://www.primelife.eu/

[24] PICOS. Privacy and identity management for community services. URL http://www.picos-project.eu/

[25] EnCoRe. EnCore: Ensuring consent and revocation. URL http://www.encore-project.info/

Several examples exist where new technologies that could be considered harmful have been modified without compromising the goals of the application. For example: many items in shops are now tagged with RFID labels, a technology which has been criticised because labels can be read from a distance of 6-9 metres allowing an eavesdropper to scan your shopping after you have left the shop. The IBM 'clipped tag'[26] is essentially an ordinary RFID tag but has been designed so that the customer can rip the antenna from the RFID label after purchase; the tag can still be read, if say, the item is returned, but only at a distance of 5 centimetres. In the UK, Marks & Spencer have integrated RFID tags into the barcode labels on many of the items of clothing in their stores. There are no RFID scanners at the store tills implying that the information in the RFID label will not be associated with the purchaser. The labels can be removed from the item before leaving the store.

Similarly, at the University of Toronto, Canada a privacy-enhancing approach to video surveillance is being developed. They have found a way of removing the personally-identifying parts of an image such as the face or body, and storing these separately from the rest of the image in a secure way. The footage can be scanned for suspicious events and if an incident merits further investigation, such as a crime scene, the police can decrypt the video content in order to identify the subjects[27].

Electronic voting, (e-voting), deserves to be mentioned in this section. Many governments are keen to realise the potential benefits such as increased voter participation but most electronic voting systems employed around the world today are not verifiable. In the UK in 2007 a series of national, e-voting pilots were conducted; these were criticised because the software and technology used could not be guaranteed to fulfill the basic requirement of verifying that the votes cast were counted as cast. E-voting has been an active field of research for the last 30 years. In the UK researchers at the University of Surrey and the University of Newcastle are working on Prêt à Voter (PAV)[28]. PAV has a touchscreen interface and uses paper based ballot forms that are turned into encrypted receipts providing both security and auditability. It is resistant to attempts to coerce users to vote in a particular way and easy to use. In simple terms PAV does not trust any of the software or equipment used in the voting and counting process but allows a voter to check every aspect of the election.

Researchers in the UK and Finland have been finding out what your laptop is saying about you[29]. There is a large amount of information leaked from wireless-enabled laptop computers, from the moment you switch them on, to initialise wireless network connections, shared file systems and printers and other services. Worse still, this information is leaked on a regular basis as the operating system will retry to connect periodically. The details revealed may enable identification of a user's corporate affiliation, user name, email address and even their real name. A partial solution to the "chattering laptops" problem does exist: some operating systems have the capability to detect automatically whether, say, the laptop is in the office and will enable or disable system services accordingly.

---

[26] Paul A. Moskowitz, Andris Lauris, and Stephen S. Morris. A privacy-enhancing radio frequency identification tag: Implementation of the clipped tag. *Pervasive Computing and Communications Workshops, IEEE International Conference on*, 0:348–351, 2007. http://tinyurl.com/5m386v

[27] Ann Cavoukian, Ph.D, Information and Privacy Commissioner, Ontario, Canada. Privacy and radical pragmatism: change the paradigm. Technical report, 08 2008. URL http://tinyurl.com/6mqlww

[28] Prêt à voter - verifiable electronic elections. URL http://www.pretavoter.com/

[29] Tuomas Aura, Janne Lindqvist, Michael Roe, and Anish Mohammed. Chattering laptops. In Nikita Borisov and Ian Goldberg, editors, *Privacy Enhancing Technologies, 8th International Symposium, PETS 2008, Leuven, Belgium, July 23-25, 2008, Proceedings*, volume 5134 of *Lecture Notes in Computer Science*, pages 167–186. Springer, 2008. ISBN 978-3-540-70629-8

Future challenges and trends

In this section we have invited privacy experts to comment on future challenges and trends. Technological advances are increasing apace and there is no doubt that the opportunity for collecting and collating data about us will also increase proportionately, we need only consider that some cars already contain over 60 microprocessors. The term pervasive computing refers to a vision of the world in which small, networked devices are seamlessly embedded in our environment and collaborate without our conscious intervention. For example: clothing sensors that monitor body temperature could be used to control heating and lighting system or a sensor attached to an Olympic athlete's body might relay readings on speed, stride frequency and stride length to the team coach. Researchers are working on these technologies with the aim of producing commercially available products in the near term.

Professor Stephen Hailes, University College London, UK, advises that we remain vigilant of the privacy implications:

"My feelings are that pervasive computing technologies potentially allow a level of intrusion into the lives of individuals far greater than ever before possible. Moreover, such devices are purposely built to be invisible, and are designed so as to be sufficiently cheap that they can pervade many of the aspects of our lives. The design aims make the technology extremely useful - it is capable of providing assistance in much of our lives and the non-networked versions already do; in particular it is capable of adapting to our needs as a consequence of the information we supply it with or that it can learn. Consequently, precisely the same attributes that make the technology useful also make it potentially rather dangerous if left completely uncontrolled."

PETs can provide a way of harnessing new technologies, allowing us to accrue the associated benefits without undermining our rights to privacy. Dr Ann Cavoukian, Ontario's Information and Privacy Commissioner, believes that using PETs as transformative technologies represents a way forward:

"By applying a privacy-enhancing technology to a surveillance technology, in a positive-sum paradigm, you develop what I am calling 'transformative technology' - transformative because you can, in effect, transform the privacy-invasive features of a given technology, rendering it privacy-protective … the effect is to minimise the unnecessary collection and use of personal data by the system, while strengthening data security - win/win, not either/or."

As ever, it is important that we think about the privacy implications of new systems and technologies from the outset. Dr Ian Brown, Research fellow at the Oxford Internet Institute, Oxford University, maintains that privacy-preserving techniques cannot be considered as an afterthought:

"We now know how to design systems that protect privacy while maintaining security. We just need to see the political will to make privacy by design a fundamental part of the information system and policy development processes of government and commercial organisations."

The drive towards user-centric identity management will continue to be an important way of protecting the individual online, Caspar Bowden, Chief Privacy Adviser, EMEA, Microsoft commented:

"...the kind of specialised cryptography in U-Prove (and IBM's IDEMIX) which allows one to do 'authentication without identification' is absolutely pivotal"

Caspar would like to see the technology exploited in a much wider range of scenarios in which the collection of personal data is reduced to an absolute minimum. Potential applications include the ability to prove age or other personal attributes without revealing any additional personal information. Similarly, it is possible to prove entitlement without revealing identity. In applications such as: privacy-protecting road pricing, congestion charging schemes, receipt of welfare benefits or health care the use of these technologies would ensure the secure transfer of appropriate, authenticated information thus reducing the threats to privacy or fraud.

For the enterprise, the development of privacy-enhancing architectures offers a promising way forward. Dr Stuart Shapiro, Principal Information Privacy and Security Engineer, The MITRE Corporation:

"We are starting to see increasing numbers of enterprise PETs aimed at helping large organisations (data stewards) manage more effectively and responsibly the personally identifiable information (PII) they collect and use. I anticipate that this trend will continue, especially as current technologies prove their value. I also expect to see continued development of PETs aimed at individuals (data subjects), perhaps with more nuanced objectives than strictly preventing the collection of PII in the first place. Indeed, we may at some point see enterprise and/or personal PETs integrated with infrastructural services such as identity management. Such integration could help move us away from PETs that target specific problems and toward a more architectural perspective in which privacy is viewed as a general system property rather than as a set of discrete controls."

About the Enterprise Privacy Group

The Enterprise Privacy Group (EPG) is an independent think-tank and membership body dedicated to creating innovative privacy management solutions. Private companies and government departments join the Group to develop solutions for their privacy, data protection and freedom of information concerns. In return for a single annual subscription, organisations have free access to a broad range of workshops, training courses, practical research and an annual conference.

Membership

EPG offers a membership body that is open to commercial, government and academic organisations from around the world. EPG Members benefit from meetings, workshops, reports, training courses, conferences and expert advice that will help them to minimise privacy-related risks in a highly cost-effective way.

Professional Services

The Enterprise Privacy Group also offers a broad spectrum of privacy services, including risk assessment, policy development, process implementation, training and awareness, data protection audits and recruitment of privacy officers on behalf of clients. EPG's team of experienced privacy professionals offer a unique, independent and ethical service for organisations that wish to set the highest possible standards for the handling of personal information.

Independence

The views expressed in this document do not necessarily reflect those of the Group's Member organisations.

The Enterprise Privacy Group is sponsored by:

BAE SYSTEMS

hp invent

csia

Microsoft

ORACLE

Sun microsystems

symantec

TATA CONSULTANCY SERVICES

T-Systems

Enterprise Privacy Group
Old Bank House
59 High Street
Odiham
Hants RG29 1LF
T: 01256 702325
www.privacygroup.org