Cyber Security
Knowledge Transfer Network

Knowledge Transfer Networks
Accelerating business innovation;
a Technology Strategy Board Programme



Privacy Engineering Whitepaper

A Report from a Special Interest Group of the Cyber Security KTN

Dr Steve Marsh
Dr Ian Brown, Oxford University
Mr Fayaz Khaki, Ernst & Young

Executive Summary

Privacy is a hot topic – one that enterprises ignore at their peril. A stronger legal and regulatory environment, high profile privacy failures, and increasing public concerns build the case for enterprises to take privacy seriously.
For those new to the subject, this paper describes the harms that privacy failures can lead to, and the reasons why privacy issues must be addressed. Harm may happen to individuals, to organisations, or to society as a whole, and enterprises should address the effects on all of these when contemplating new information systems. Leadership is essential if concern for privacy is to be embedded throughout an organisation's culture, processes and systems.

For those attempting to design privacy in to their systems, this paper provides guidance on the issues that must be addressed. The range of issues is broad, and we can only scratch the surface here. More work is needed to develop the detail, and we hope this paper will inspire that development. But the breadth and complexity of the issues also emphasises the need to develop skills and ethics within a profession of privacy practitioners.

Finally, this paper offers three clear conclusions about the nature of privacy issues, who is responsible, and how the threat of breaches can be vastly reduced by taking swift and appropriate measures.

We welcome interested parties to join the Privacy Special Interest Group of the Cyber Security Knowledge Transfer Network, to help take these ideas forward.

1. Introduction

Privacy is a hot topic. New social and business models are springing up as the result of web technologies and the plummeting cost of information storage and processing. There are now many high profile cases where features or failures in information systems have crossed the line of the public's, often un-stated, expectations of privacy – causing outrage or distress to users, and reputational and financial loss to organisations. But there is little guidance available to those who commission, design, build or operate information systems to help them avoid such disasters.

Many information systems are designed to collect and store all available data, because filtering and selecting takes more effort, and the benefits of investment in the system may grow if new ideas to extract value from data emerge in the future. The value of data itself is thought to rise, as more data is aggregated.

But organisations that adopt this approach are doing more than storing up information – they are also storing up problems for themselves as public concerns over privacy grow, privacy issues rise up the political agenda, and legal sanctions for privacy violations increase. The cost of managing this data is often not considered, but can become very significant as disclosure or freedom of information requirements grow.

This Whitepaper sets out the case for organisations to address privacy issues comprehensively, and begins to provide practical advice to organisations on how to do this. The subject really requires several books to cover the topic

comprehensively, and clearly a single whitepaper can do no more than scratch the surface. However, the authors hope that this paper will encourage others to develop the themes.

The paper starts by clarifying what we mean by privacy, then addresses the 'business case' for organisations to address privacy issues, and sets out a range of measures that organisations should consider when implementing information systems.

This work was carried out by a Special Interest Group of the Cybersecurity Knowledge Transfer Network, and was initiated following a 'Fine Balance' Symposium held in December 2006 that was organised jointly by the Cybersecurity, Location & Timing, and Sensors KTNs.

DISCLAIMER

The views expressed herein are entirely those of the authors, and do not represent the organisations to which they are affiliated.

Dr Steve Marsh
Dr Ian Brown, Oxford University
Mr. Fayaz Khaki, Ernst & Young

The authors acknowledge and appreciate the discussions held with fellow members of the Privacy SIG.

2. Privacy – What is it?

2.1 A Human Right

The United Nations Universal Declaration of Human Rights[1] and the European Convention on Human Rights[2] both enshrine the right to respect for private and family life, home and correspondence. The UK has adopted freedom from arbitrary interference with privacy into law.[3]
The UK Data Protection Act[4] implements the European Data Protection Directive. [5]

2.2 The Royal Academy of Engineering report

The Royal Academy of Engineering published a report on 'Dilemmas of Privacy and Surveillance' in March 2007[6] in which they highlighted the many and varied forms of privacy, for example:

[1] http://www.un.org/Overview/rights.html Article 12.
[2] http://www.hri.org/docs/ECHR50.html Article 8.
[3] http://www.opsi.gov.uk/ACTS/acts1998/ukpga_19980042_en_1
[4] http://www.legislation.gov.uk/acts/acts1998/ukpga_19980029_en_1
[5] http://ec.europa.eu/justice_home/fsj/privacy/index_en.htm

- As confidentiality of information about ourselves or things we do;
- As anonymity, preventing the association of information about ourselves or things we do being associated with real identities;
- As protection of identity, such as the separation of an official role from a personal identity;
- As the right to make certain choices without interference from others;
- As the right to be left alone;
- As control of information about ourselves.

The 'Dilemmas' in the title of the report reflects the fact that none of these rights can be absolute. We ourselves, or other people (often reflected in 'society' as a whole), may reasonably expect some of these rights to be encroached upon in return for services offered, or to deter or detect crime, or promote personal accountability in a relationship.

2.3 A taxonomy of privacy

'Protection of privacy' in a system therefore depends on reaching an appropriate balance of risk and benefit. Complicating factors are that risks and benefits will be different for individuals affected by the system, the organisation owning the system, and society as a whole; and that perceptions of, and impacts on, privacy are changing rapidly.

Prof. Daniel J. Solove of the George Washington University, has developed a helpful 'taxonomy of privacy'[7] by focussing on the 'harms' that arise from privacy violations, rather than on the more abstract concept of privacy itself. This taxonomy is a helpful basis on which to develop a risk/benefit analysis. Prof. Solove classifies harms under the categories of Information Collection, Information Processing, Information Dissemination, and Invasions. The following draws heavily on his work.

2.3.1 Information Collection

Harms classified under 'Information Collection' are 'surveillance' and 'interrogation'.

Surveillance is the passive form of information collection, and may be carried out in public or private spaces. The boundary between public and private is also blurring with, for example, the use of web browsing tracker technologies. Harms range from the oppressive effect on individuals of 'being watched' – often cited as a reason for moving from a village to a town – to the 'chilling effect on society' where pervasive surveillance can curtail the free association and espousal of radical ideas that are essential to a democratic society.

6

http://www.raeng.org.uk/policy/reports/pdf/dilemmas_of_privacy_and_surveillance_report.pdf

7 http://papers.ssrn.com/sol3/papers.cfm?abstract_id=667622

Interrogation is the active form of information collection.
Harms range from the offence felt by probing questions that might elicit a response of 'Mind your own business', to the distortions and forced betrayals that were seen, for example, in the McCarthy era.

Even individuals with 'nothing to hide' may react strongly against information collection. The UK's Department for Transport found this out a few years ago in a research project on motorway usage: using information from an automatic number plate reader and the database of registered vehicle keepers, they sent questionnaires out asking for information on the reasons that the observed journeys had been made. Many respondents were outraged at this intrusion of privacy, despite the journeys having been in a public place, and the existence of both the number plate reading technology and the database of registered keepers being well known.

2.3.2 Information Processing

Harms classified under 'information processing' are 'aggregation', 'identification', 'insecurity', 'secondary use' and 'exclusion'.

Aggregation is the process whereby different items of information about a person are pulled together to infer new 'facts', and is increasingly prevalent with advances in database technology.

Harms can range from the unease felt by an individual when someone else knows something about them that the individual has not consciously released, through to discrimination against an individual based, for example, on their purchasing power or genetic make up.

While aggregation of information can occur with respect to a pseudonymous individual, *identification* is the process of associating information with a particular human being and may or may not involve information aggregation. Identification can lead to harm to an individual in conjunction with insecurity or disclosure if, for example, the individual has expressed views that others violently disagree with (e.g. the murder of the Dutch film maker Theo van Gogh in 2004) or is in a witness protection programme. Groups of individuals can also be harmed if identification permits discrimination, such as happened to the Jews in Nazi Germany.

Insecurity of collected information may involve illicit aggregation, modification, or use, or unauthorised access and dissemination.
Harms can range from distress and uncertainty caused by knowledge that a security breach may occur, to theft of credit card information or identity credentials. While some view theft of credit card information as relatively benign (since banks will usually compensate the individual if they have not been negligent) such theft has a significant cost to society as a whole, and if the credit card information is used to commit an illegal act then there can be serious implications for the individual.

Secondary use is the use of data for purposes unrelated to the purposes for which the data was collected, without the individual's consent. As the cost of storing and processing information continues to decrease and the understanding of the value of information assets increases, the temptation to engage in secondary use becomes increasingly great.

Prohibition of secondary use is one of the Data Protection principles set out in the UK Data Protection Act. Harm arises because individuals' expectations of the use to which their data is put are not met, and this can damage the trust relationship between the individual and the information processor. If secondary use becomes widespread, individuals may become increasingly reluctant to divulge information in the first place.

Another key data protection principle is that of subject access rights – the right to see, and if necessary, correct, information held about the individual. *Exclusion* is the harm arising from failure to provide this access.

The harm arises because of the powerlessness and frustration that an individual may experience if such access and correction mechanisms are withheld, and because this withholding in itself may decrease the accuracy of information used leading to other harms.

2.3.3 Information Dissemination

Harms classified under 'information dissemination' are 'breach of confidentiality', 'disclosure', 'exposure', 'increased accessibility', 'blackmail', 'appropriation' and 'distortion'.

Breach of Confidentiality occurs when information is disclosed in a situation where the individual particularly expected a confidence to be kept. The harm caused differs from the disclosure of the information itself (considered in the next paragraph) since it is the damage to the trust relationship between the individual and the disseminator that is significant, rather than any damage caused by others having access to the disclosed information. Such cases usually arise where there is a significant expectation of professional confidentiality – for example, between a patient and a doctor, or a client and a solicitor.

Disclosure occurs when true information about a person is disclosed to others.

Harms can range from publication of paparazzi photographs, publication of private information that – whether or not it could lead to reputational damage – is not in the public interest, through to publication of information (as in the case of identification) that could lead to physical harm. Solove points out that publication of true, but partial, information can lead to a distorted assessment of a person, and can also lead to people becoming 'prisoners of their past' – unable to redeem themselves from past errors of thought or deed.

Exposure is the exposing to others of certain physical and emotional attributes of a person. Solove describes 'grief, suffering, trauma, injury, nudity, sex, urination and defecation' as involving 'physical, instinctual, and necessary' primal aspects of our lives that we have been socialized into concealing. Acceptable norms of exposure vary with time and place, but when exposure occurs the harm is against the dignity of an individual (or their family and friends) and can lead to humiliation and loss of self-esteem.

Increased accessibility arises in particular from provision of information through broadcast media and the internet. The information in question is already 'public' in some sense, but the mode of presentation makes it available to a much wider range of people than originally intended.

Harm can arise as the information can be taken out of context, given permanency, used for purposes for which it was not intended, and used very easily by people who would not otherwise go to the trouble of collecting the information. Recent examples are the posting on YouTube of videos taken by pupils of their teachers, or self-posting by students to social networking sites such as MySpace, of accounts of their drunken exploits. These examples have led to humiliation, harassment and denial of employment.

Blackmail has a long history of illegality, irrespective of whether the information used for the blackmail was true or false, and irrespective of whether the victim was actually committing a crime for which they were being blackmailed.

Society judges the harm from blackmail to arise both from the nature of the act itself – in a similar fashion to extortion, kidnapping and protection racketeering – and also because of the coercion of the victim, even if the victim might appear to consent to the deal.

Appropriation occurs when one's identity or personality is used without consent for the purposes and goals of another person. This may range from using one's image to endorse a product or opinion, to theft of a complete identity.

Harm arises through the false attribution of beliefs or actions to an individual, which may lead to humiliation, reputational damage, financial loss or, if the appropriation is used to commit further crimes, the threat of incarceration.

Distortion is the manipulation of the way a person is perceived or judged by others. Solove states that this consists of being exposed to the public inaccurately, by the publication of false and misleading information, or selective revelation of the truth.
The harm is caused not just to the individual who's reputation is damaged, but also to the relationship that the rest of society may have with that individual.

2.3.4 Invasions

Harms classified under 'Invasions' are 'intrusion' and decisional interference'.

Intrusions are invasions or incursions into one's private life. They involve the interruption of one's activities by the unwanted presence or activities of another.

Historically, intrusions have been physical – into one's house or personal space. But the introduction of communications media have also brought intrusion through spam, junk mail, telemarketing and so on. The harms identified are disruption, interference with seclusion, and interference with association and ideas.

Finally, *decisional interference* is the harm caused by the interference by others – particularly governments – in decisions that should primarily be an individual's.

Examples range from the books one chooses to read, to the medical treatments one chooses to accept.

Harms can range from undermining of personal responsibility to chilling effects on society similar to pervasive surveillance.

2.4 Consequences

A major step in privacy engineering should be the completion of a Privacy Impact Assessment. This is discussed in more detail in a later section of the report. But it is clear from this description of the potential harms arising from privacy violations that this impact assessment should consider the potential impact not just on the organisation, but also on individual 'users' of the system and society as a whole; and that this assessment should consider not just the impact if the system works as expected, but the potential impact should the system fail.

'Counter-measures' introduced to protect privacy should also be assessed for their ability to address each of the potential harms described above.

3. Privacy – Why respect it?

Privacy is still a relatively little known issue for enterprises, and compelling arguments may be needed to ensure that it gets on to the agenda of busy executives. Such arguments are still immature, but are strengthening as more real-life incidents accumulate. Depending upon the type of enterprise, one or more of the following may resonate.

3.1 Legal Obligations

Organisations are required to consider these questions across the European Union under the Data Protection Directive, and in the many other countries around the world that have introduced compatible legislation (including under the "Safe Harbor" negotiated between the European Commission and the US Federal Trade Commission). The Article 29 Working Party set up under the Data Protection Directive has published a number of documents that explain to businesses their legal responsibilities.[8] There are also international standards on privacy controls, including ISO7799.005 12.1 / 15.1.4.

3.2 Liability

Customers whose privacy rights are violated by an enterprise may be able to sue under a range of privacy laws. TJX, the company which, in 2007, disclosed the theft of 44 million credit card details, has reportedly set aside $107m to settle a class action suit brought against it by customers in the U.S., Canada and Puerto Rico.

3.3 Reputation

Even if no direct financial liability applies, an enterprise may suffer a damaging loss of reputation if it ignores its customers' privacy. In November 2007, the social networking site, Facebook, had to back down over the introduction of its 'Beacon' feature, which automatically notified the 'friends' of any user making a web purchase at participating sites of the nature of the purchase. Despite claims that Facebook users have a different attitude to privacy, the outcry that this feature led to, caused Facebook first, to add an opt-out option on a per site basis, and then to add a global opt-out facility. This led to further questions about privacy of information on Facebook, including enquiries from the Information Commissioner about the effectiveness of deletion of profile data following users' requests.

3.4 Information Access and discovery

Legal compliance is a key area that will determine the lifetime of some items of personal information held by organisations. In the case of credit card transactions, for example, fraud complaints can be made up to 6 months after a payment is made – meaning that organisations will need to keep related

[8] See http://ec.europa.eu/justice_home/fsj/privacy/workinggroup/wpdocs/2007_en.htm

data for that period. There has been recent controversy in the US where stores have complained about the security and privacy risks that result from the Payment Card Industry Standards Council's requirements for them to retain certain credit card data.[9]

A very small number of businesses such as public telecommunications providers are required under specific crime-related legislation (e.g. the Data Retention Directive) to store personal information for non business-related purposes.

3.5 Corporate Social Responsibility

In contrast to risks to enterprises noted above, privacy may reach the Board agenda as an example of corporate social responsibility that builds the image and reputation of an organisation. Just as, for example, environmental issues are now commonly given prominence in the motor industry, so privacy could become a factor for the 'cyber-environment'.

3.6 Example incidents

2007 saw a number of high-profile privacy breaches in addition to those mentioned above, that illustrate areas that enterprises should consider.

3.6.1 H.M. Revenue and Customs

On 20th November 2007, the Chancellor Alistair Darling made an emergency statement in Parliament about the loss of two discs containing details of 7.25 million families in the UK who had children under the age of 16. The discs disappeared in transit via courier to the National Audit Office. The Permanent Secretary of HMRC resigned. [10]

3.6.2 Jeremy Clarkson

Following the HMRC incident, the television presenter Jeremy Clarkson claimed in his column in the Sun newspaper that "…he'd never known such a palaver about nothing" and published his bank account details. He subsequently admitted in his Sunday Times column that someone had used these details to create a £500 Direct Debit payment to a charity. [11]

3.6.3 Norwich Union

On 17th December 2007, the Financial Services Authority fined Norwich Union Life £1.26m for exposing its customers to the risk of fraud, through lax security controls at its call centres. [12]

---

[9] Jaikumar Vijayan, Retail group takes a swipe at PCI, puts card companies 'on notice', Computerworld, 4 October 2007

[10] http://news.bbc.co.uk/go/pr/fr/-/1/hi/uk_politics/7103566.stm
[11] http://news.bbc.co.uk/go/pr/fr/-/hi/entertainment/7174760.stm
[12] http://www.fsa.gov.uk/pages/Library/Communication/PR/2007/130.shtml

3.6.4 Department of Health

On 19th December 2007, the Information Commissioner's Office found the Department of Health in breach of the Data Protection Act, following an alert in May that sensitive personal details of junior doctors were accessible to anyone accessing the website of the Medical Training Application Service. [13]

3.6.5 Land Registry

In November 2007, the Land Registry removed scanned mortgage deeds and leases from its "Land Register Online" web site, following a number of cases of fraud. Land Registry had paid out just under £12m in fraud and forgery claims in 2005/06 and 2006/07. [14]

[13] http://www.ico.gov.uk/upload/documents/pressreleases/2007/doh_undertaking_pr.pdf
[14] http://news.bbc.co.uk/go/pr/fr/-/1/hi/uk_politics/7079494.stm

4. Privacy – How to design it in

It has long been known that security can not be successfully 'bolted on at the end' of a project. This fact is even more applicable in the case of privacy. An organisation must commit to protecting its customers' privacy up front; that commitment must imbue the culture and ethos of an organisation, and must shape the approach to projects even before design starts.

Privacy requirements must be fed in at each of the four stages – initiation, planning, execution and closure – of a generic project lifecycle. The requirements and level of detail must align with the objectives of each output of the project lifecycle stage to ensure timely information is made available to the project.

At *Project Initiation* high level privacy objectives need to be set. At the end of the project initiation stage, the enterprise should have a general idea of what the project will entail and what assets will be involved, introduced or addressed by the project. Therefore the privacy requirements at this stage should involve high level recommendations. For example, the project owner will need to be aware of applicable privacy laws and regulations, such as the EU Data Protection Directive or the US Safe Harbour agreement. Technology envisaged for use by the project should also be subject to a high level review to ensure that appropriate privacy controls can be implemented. As a result of this, the project will be aware of, and be able to factor in, the appropriate privacy requirements at the next lifecycle stage.

The *Project Planning* stage in the lifecycle enables the enterprise to develop detailed requirements for the project. The privacy requirements at this stage will need to be just as definitive. They will enable the enterprise to take account of applicable laws and regulations. Further privacy requirements relating to information assets will be made clear. Technology that will be introduced as a result of the project will also be defined. The privacy requirements for the technology will be appropriately and accurately described. Mechanisms such as encryption will be proposed to protect, for example, consumer and client data on storage media. Additional privacy requirements will be made if necessary, for example to ensure that identifiable data is not made available when viewed by personnel who do not have a "need to know".

During the *Project Execution* phase, any problems relating to the privacy solutions proposed for the project should be identified, and decisions made on alternative solutions to ensure that the required level of privacy is still maintained. The project should document any privacy related issues and exposures. The project owner should be made aware of these prior to the project 'going live' in the production environment, and should sign-off any privacy related issues that have been raised by the project.

Privacy requirements need to be consistently and continually addressed after *Project Closure*, in the production environment. This should be achieved by way of audits and change control procedures. Internal audits should be conducted at regular intervals to ensure that there aren't any breaches of the agreed privacy implementations and policies.

Finally, privacy must also be addressed at *System Decommissioning* – whether this involves secure deletion of data from computer media before disposal, or destruction of paper records before vacating buildings, for example.

The remainder of this section sets out some general principles of privacy and data protection that should shape enterprises and their thinking about projects; some aspects pertinent to organisation and culture; and a range of measures to apply throughout the project lifecycle.

4.1 Principles

Clear principles are needed to underpin the design and implementation of an information system. The UK Data Protection Act sets out 'Data Protection Principles' that remain very valid in this context, and are set out again here with the addition of others that address user preferences.

4.1.1 Fair and Lawful

Information must, of course, be processed in accord with the law in the relevant jurisdictions. But the first data protection principle also states that information should be processed 'fairly'. This is an important aspect that can enhance or degrade the reputation of an organisation.

4.1.2 Defined scope

The second data protection principle states that personal information shall be obtained only for one or more specified and lawful purposes, and shall not be further processed in any manner incompatible with that purpose(s). The scope of an information system must be well defined at the outset, and 'scope-creep' eliminated.

4.1.3 Minimization

The third data protection principle states that the collection of information shall be adequate, relevant and not excessive in relation to the purpose or purposes for which it is processed. It is important to avoid collecting excessive information, both because this causes additional information management problems for the organisation, and also as it increases the privacy challenges. Anonymizing personal information, or permitting pseudonyms, can be an important privacy enhancing technique, though limitations of this approach need to be understood.

4.1.4 Accurate

The fourth data protection principle states that stored personal information shall be accurate and, where necessary, kept up to date. There is a clear business advantage in complying with this principle, and this can be a critical factor for the individual, for example, in a healthcare context.

4.1.5 Time limited

The fifth data protection principle states that information processed for any purpose or purposes shall not be kept for longer than is necessary for that

purpose or those purposes. This is an increasingly important principle for organisations as their information management challenges – for example, for compliance or legal discovery - can grow enormously if appropriate data retention policies are not implemented. It is also of importance to the individual to control their biographical footprint.

4.1.6 Authorised processing

The sixth data protection principle states that appropriate technical and organisational measures shall be taken against unauthorised or unlawful processing of personal data and against accidental loss or destruction of, or damage to, personal data. This principle emphasises the importance of good information security in support privacy – though privacy itself is more than just information security.

4.1.7 Consent

The data protection principles apply independently of whether or not the individual concerned has given consent to the processing of their information. Nevertheless, obtaining the consent of the individual is a valuable safeguard both for the individual and for the organisation. Particular attention should be paid to ensuring that consent is *informed,* and the difficulty of achieving, and demonstrating this, should not be under-estimated. An additional important principle and protection for the individual is that they should be able to retract their consent at some future date, which should lead to removal of their personal information form the system (subject to legal retention requirements.)

4.1.8 Choice

Customers should also as far as possible be given the choice of opting in to services that require the collection of additional personal information. For instance, an online book store might allow customers to choose to shop anonymously, deleting delivery addresses as soon as an item is dispatched and keeping payment information only as long as required by the relevant payment system. Alternatively, customers might choose to have their address saved to make future purchases easier – and for their list of past purchases to be stored and data mined to improve the quality of marketing messages presented on future visits. Again, an important element is that the customer may alter their choice at some future date, with consequent deletion of records from the system.

4.1.9 Responsibility

'Privity' is a long established legal principle that may be relevant in this field – that the customer need only take action against the final link in a supply chain – legal remedies then apply up the chain to the failing party.

Where organisations are at the head of information supply chains they will need to ensure compliance of other organisations in the chain to their privacy protection policies.

4.1.10 Transparency

Organisations should ensure *Transparency* for individuals with a clear and available description of the data that is held about them, the purposes for which it is used, the people who have access to it, and the mechanisms for remediation.

4.2 Organisation and Culture

4.2.1 Ethics

The ethical stance of an organisation will have a large effect on the ease with which privacy controls are designed in to a system. Just as organisations adopt ethical positions on environmental impact, or fair trade, so should they adopt a clear position on privacy – building reporting on this position into Board reporting through mechanisms such as the Balanced Scorecard.

4.2.2 Privacy Officer

The Data Protection Act defines the role of data controller, but in practice this has been devolved to relatively junior positions in the organisation. Organisations should ensure that a senior role is established with *overall responsibility* for privacy.

4.2.3 Information Management

Organisations with clearly documented business and information management processes will be at a great advantage in enforcing controls on the processing of personal data, as well as implementing requirements of data protection law such as the ability to answer subject access requests. They will also be much better placed to understand the value to their business of different pieces of customer information.

4.2.4 Staff awareness and training

As in other walks of life where significant impact on the public can occur, there may be a need for a professional privacy body, and formal qualifications including degree accreditation. Organisations will, at least, need to make staff aware of the expected approach to customer privacy, and ensure that this awareness is maintained.

4.3 Programme Definition and project commissioning

4.3.1 Governance

Clear lines of responsibility are vital for successful programme delivery. Each programme or project should have a designated Senior Responsible Owner who is accountable and responsible for risk and privacy management and benefit delivery.

4.3.2 Programme scope

A well defined scope for the programme or project is both a prerequisite for successful delivery, and an important starting point for privacy engineering.

Its absence violates the second data protection principle, and can lead to scope creep.

4.3.3 Privacy Objectives

The Senior Responsible Owner should set privacy objectives at an early stage of the programme, in a similar way that security targets are set. These objectives then inform subsequent design decisions. As with security, it is difficult, if not impossible, to add-in privacy late on in a project.

4.3.4 Independent Oversight

Organisations should consider mechanisms to deliver independent oversight of privacy arrangements, to build trust amongst their customers and clients.

4.3.5 Breach notification

Organisations should consider whether and how they should implement procedures for notifying individuals of privacy breaches. This is a legal requirement in an increasing number of U.S. states, and is being considered in certain industry sectors in Europe.

4.4 System design

4.4.1 Architecture

Privacy objectives need to be addressed right at the start of the architectural design process. Decisions made at this stage can have a major impact on the cost and difficulty of achieving these objectives. For example, if the 'minimization' principle is not applied strongly enough in designing data collection, much more effort may need to go into privacy protection measures during data processing and storage.

The use of 'legacy' systems needs also to be carefully addressed, as these may have been designed and built before the organisation's current privacy policies were in place.

4.4.2 Privacy Impact Assessment

A Privacy Impact Assessment should be started early during the design phase, and the output from the PIA iterated with re-design until the Privacy Objectives are satisfied. The assessment should not just address the privacy outcomes when the system is working correctly. Rather, it should also address the harms accruing to the organisation, the customers and to society more generally arising from potential privacy violations caused by possible failure modes of the protections built into the system. Failures will occur in any system – the design of the system should be such as to limit the damage when something goes wrong, allow early detection, and permit easy repair.

4.4.3 Information Security

Good information security practice should be adopted, including physical and personnel security measures, audit and data management, configuration control, intrusion detection and incident response. Organisations should implement appropriate information security safeguards to deter, detect and

prevent misuse and abuse of personal information. Independent accreditation against security standards such as the ISO27000 series should be considered.

4.4.4 Access control

Systems should implement strong access controls, to ensure that personal information is only accessed by those who are authorised to do so. Accesses should be accounted, and logs regularly audited.

The principles of 'least privilege' and 'separation of authority' should be applied. For example, the security administrator should have access to audit functions but no access to system administration functions. The system administrators should not have access to audit functions or customer personal information.

Organisations should also consider *'Intermediation'* as a means of implementing separation of authority and increasing the transparency of information use. People might lodge their information with trusted parties who control access to that information on the individuals' behalf. The Royal Academy of Engineering report contrasts these 'Big Brother' and 'Little Sister' models.

4.4.5 Metadata & permissions

Where possible, personal information should be stored together with metadata that describes it and its intended use. In conjunction with efficient access control, security and privacy policy can be enforced through application of granular permissions.

4.4.6 Consent & Retraction

Organisations should implement transparent procedures for giving and retracting consent by individuals for the processing and storage of personal information, and for expressing (and changing) choices, wherever possible, for the personal information held.

4.4.7 Remediation

Organisations should implement transparent procedures for remediation of errors in personal information, or privacy breaches.

4.4.8 Redress

Organisations should implement transparent procedures for redress for customers and clients in the event of errors in processing personal information or privacy breaches.

4.4.9 Development and test

Real customer data may need to be used during system development and test. Particular care should be taken to protect privacy during these phases as the full range of privacy and security controls are unlikely to be in place. In particular, if development is outsourced, the organisation should be diligent in ensuring that the organisations to which the work is outsourced understand and comply with the privacy and security policy.

4.5 System implementation

4.5.1 Encryption

The use of data encryption is a key mechanism to enforce access controls against those with physical access to a transmission or storage medium containing personal data. Driven by the Data Protection Directive in Europe and security breach notification laws in the US, it is also rapidly becoming an essential tool for the protection of personal data.

Encryption can be used to efficiently "delete" large amounts of personal information, which may be stored in several different locations (such as backup sites) and be difficult to definitively wipe. If the data is encrypted, deletion of a carefully controlled associated decryption key effectively prevents further access to that information. Automatic deletion of keys on a pre-programmed expiry date can thus "age" and permanently wipe any quantity of data. Trusted computing hardware can be used to manage critical cryptographic keys.

4.5.2 Digital Rights Management

The application of encryption at document or record level is being integrated into broadly based digital rights management (DRM) systems. These systems can be used to enhance privacy by controlling the actions that a recipient of a document can undertake: for example, view but not print; only view from a designated machine; or only access before a specified date. Hardware support for DRM is being implemented in mass-market PCs via the 'Trusted Computing Platform.' Features in the platform should be considered for use in enhancing privacy and security of endpoints.

4.5.3 Mandatory Access Controls

Organisations should also consider designing systems to tag items of personal data with information on its sensitivity, permitted uses and lifetime (along the lines of Mandatory Access Controls in operating systems such as Security Enhanced Linux.)

4.5.4 Privacy Enhancing Technologies

There is continuing research in the field of privacy enhancing technologies. These often use encryption in novel ways to provide strong privacy properties. No one technique has achieved widespread take-up, but blind- and partially-blind signatures and the associated concept of limited-show certificates appear to have great potential, although conceptually difficult to grasp. Organisations should consider whether such techniques could contribute significantly to their privacy objectives.

4.5.5 Legal compliance

In some sectors, there is a need to retain personal or other information to comply with legal requirements. Organisations should consider the use of public-key cryptography to 'seal' such stores so that they can only subsequently be accessed by authorized entities.

4.5.6 Secure deletion

Organisations should securely delete personal information when it is no longer required. Magnetic media may need to be explicitly over written, as 'deletion' of a document from a file system does not usually guarantee erasure of the underlying data. Provision should be made for physical destruction of optical media, or paper based records.

4.6 System operation

4.6.1 Backup

All of an organisation's systems, including backup and test systems, need to be designed with minimisation requirements taken into account. If personal data is regularly deleted from main storage systems but still kept in backups it will still be vulnerable to unauthorised access. It could also be extremely costly to process for the purposes of answering subject access requests or other legal demands for information such as a court order or warrant. Appropriate security and privacy controls need to be applied at off-site backup facilities, and particular attention should be paid to the security of decryption keys.

4.6.2 Fault Repair

Access to systems is often required by third parties to undertake maintenance and repair. Privacy and security controls need to be applied during this process, and in particular, if equipment is removed from site.

4.6.3 Business continuity and disaster recovery

Particular attention should be paid to the maintenance of privacy controls during business continuity and disaster recovery activity. Shared facilities may be used for business continuity or during the disaster recovery process, and personnel may need to assume different roles at short notice. Planning and exercising is essential.

4.7 System decommissioning and disposal

4.7.1 Media

Secure deletion of data is especially important during system decommissioning and disposal, as it is at this point that the organisation's interest in the system is at its lowest. Computer equipment may be re-furbished, or appear on eBay, and some people have made some surprising finds of information from disk drives acquired through such routes. Again, physical destruction of optical and paper records is required.

5. Conclusion

This paper makes three main recommendations:

1. Organisations must take more responsibility for ensuring data privacy. The issues associated with privacy breaches are not those which can be addressed by users, they must be addressed by those planning, designing and implementing new information systems.
2. Those developing information systems and associated products/services must build privacy into every stage of the product development process. It is not sufficient to see privacy as a bolt-on to good design, integrated as a last-ditch effort to protect personal data.
3. Organisations must stop delegating responsibility for privacy issues to junior members of staff. Privacy is a board-level issue and should be the duty of a senior decision maker with the power to make important changes at the highest levels.

Mainstream interest in privacy is growing, as a result of stronger legal and regulatory regimes, increasingly frequent high-profile losses of personal data, and the realisation by organisations of the reputational, financial and legal implications of privacy failures. Enterprises must show leadership in addressing these growing risks to themselves and their clients. Given the conclusions of this paper swift action is required.

But there is little concrete support for those faced with sponsoring, designing, implementing and operating information systems, to assist in meeting privacy challenges. This paper is a starting point. The range of issues involved is broad, and emphasises the need for skills, ethics and professionalism amongst privacy practitioners.

The authors believe that much further work is needed to develop these concepts – both in the detail and as understanding of privacy develops. We hope that this white paper will stimulate that work.

The third and final ‘Fine Balance’ event - to be held late in 2008 - will look at Case Studies, particularly around the introduction of Privacy Impact Assessments.

We welcome interested parties to sign up for the Cybersecurity Knowledge Transfer Network Special Interest Group on Privacy Engineering, both to contribute to these case studies, and explore how we could take this work forward.