Framework for Information Assurance

1. The Foundation for Information Policy Research is an independent body that studies the interaction between information technology and society. Its goal is to identify technical developments with significant social impact, commission and undertake research into public policy alternatives, and promote public understanding and dialogue between technologists and policy-makers in the UK and Europe.

2. We welcome the opportunity to comment on the "e-Government framework for Information Assurance" (draft 5.1, December 2006).

3. There are four things seriously wrong with this Framework: an obsolete model of online threats, a failure to treat harm to government employees on the same basis as harm to other citizens, a failure to draw a clear distinction between identity and authority, and a security policy model that is often inappropriate.

4. This is a document that could have been largely written ten years ago, and is perfused with last-century assumptions about online dangers. What has changed the world radically since then is the advent of phishing, and of related attack techniques such as pharming. From a handful of such attacks in 2004, phishing has grown to the point that it cost UK banks eight figures, and US banks nine figures, in 2006. It is still growing rapidly and is now one of the main foci of security research and development.

5. In a typical phishing attack, victims are lured by an email to log on to a website that appears genuine but that actually steals their passwords. In pharming, the deception may occur one level deeper; the attacker somehow takes over the victims' network connections (for example, by corrupting DNS settings in a home router) and directs them to a phishing site. Deception, of various kinds, is now the greatest threat to online security.

6. The Framework teaches government systems builders to cope with higher levels of risk by increasing the intricacy of the logon procedure for example, from passwords to 'limited lifetime tokens'. But this is no longer in touch with reality. A password generator merely causes the phisherman to run a real-time man-in-the-middle attack.

7. The critical question nowadays is not how a government website can be assured of the identity of a citizen visiting it, but how the citizen can be assured of the authenticity of the website. Although there are passing references to social engineering at 7.4.1, to the need to authenticate services as well as clients at 7.5.7, and to middleperson attacks at 7.5.9, the Framework's writers show no evidence of being seized of the real problems. Designers following their advice would likely produce vulnerable, exploitable systems.

8. Many banks became particularly vulnerable to phishing because, during the dotcom boom, they largely dismantled internal controls on what customers could do online. The world's first online retail banking system (introduced in the 1980s by the Bank of Scotland) let customers only make payments to accounts that had previously been nominated in writing. A customer would write to the bank and authorise it to accept email instructions to pay their gas bill up to a maximum of X per month; a thief who obtained the customer's password could overpay their gas bill, but had no easy way of extracting cash. But once e-euphoria caused the requirement for payee nomination to

be dropped, a thief could transfer all of a customer's available balance and credit to an account overseas (or use the account for money laundering). E-banking has thus become a target for thieves, and now some banks are engaged in reimposing controls. An example at one bank is that a customer attempting to make a payment to a new payee for the first time will have the payment held up for a day, and will be sent an SMS message asking for confirmation. Such controls ought to be universal but unfortunately are not.

9. The emerging consensus of experts on phishing is that front-end authentication mechanisms simply cannot bear the full brunt of all the likely information security threats, especially to untrained users of commodity software platforms. Although it was convenient for the banks to dismantle internal controls and rely on mechanisms such as SSL, passwords and even password generators, their optimism turned out to be misplaced and expensive. The UK government seems poised to repeat their mistake, by encouraging developers to think of assurance in terms of selecting the appropriate level of authentication for an assessed level of risk. The banks' experience teaches that this just won't work.

10. The second major problem is that the Framework conspicuously shows less concern for the safety of ordinary citizens than it shows for the safety of government employees. Under the long-established rules for information classification, data whose compromise could threaten human life directly had to be classified at SECRET. This was of course mostly interpreted in the context of the human lives being those of soldiers or police officers, but as a general principle it appeared robust. Yet the Framework proposes that information whose compromise would pose a direct safety hazard to a member of the general public is only to be classified at level 3; and although level 3 is seen as equivalent to RESTRICTED, level 3 data won't actually be considered to be RESTRICTED in terms of the classification rules as that might 'cause problems' in granting subject access requests under the Data Protection Act. Introducing a new marking of PRIVATE will also further widen the incompatibility between the UK and USA protective marking systems introduced a generation ago with the UK label RESTRICTED. The added complexity, plus the autarky of a retreat still further into a national standard, may be a problem for system designers attempting to build high-assurance systems using commercial off-the-shelf components.

11. This proposed level of protective marking is extraordinarily low, both by historical standards and by comparison with practice overseas. First, the NHS's information security strategy in the mid-1990s proposed a system of security levels treating prescriptions as at a level equivalent to RESTRICTED, most medical records as CONFIDENTIAL and highly sensitive material like information on HIV/AIDS sufferers as SECRET. Second, the standard procedure in Canada is for government forms soliciting sensitive personal data to be marked 'Confidential once completed'. In addition to this understating of Infosec risks to individuals, the Framework ignores collective risks to significant numbers of people. Of course these higher levels of risk are known about in government and we do not understand why they are left out here. The overall effect is to underplay the risks faced the public, while applying much higher standards within government. Given the methods now used by terrorists and violent single-issue groups, this is just asking for trouble.

12. The third major problem with the Framework is its failure to distinguish carefully between checking identity and authority. In many applications the issue is not the identity of a user but their authority to perform some action. Names are often best bound late to credentials as this maximises modularity, facilitates delegation, protects privacy and generally minimises costs. Focussing on identity may be seen as helpful to the current Home Office's identity card agenda, but we doubt that it is prudent for departments to make their customer-facing systems needlessly expensive, clunky, and tiresome to use.

13. The fourth major problem with the Framework is that it relies on levels rather than on compartmentation – on multilevel security rather than multilateral security. CESG and other agencies and departments that have responsibility for information that the government seriously wishes to protect have been aware for many years that limiting the number of people with access to information is crucial. Aggregating large quantities of sensitive information, to which more and more people then need access in order to do their jobs, simultaneously increases the value of the target and the number of people through whose carelessness or disloyalty it can be compromised. Compartmented security policies are widely used for defence and intelligence information; and exactly the same considerations apply to information whose disclosure would harm private individuals as to information whose disclosure could harm soldiers or intelligence officers.

14. The centralisation that took place of commercial information starting a generation ago, whereby (for example) any bank teller could enquire about any customer's account, certainly led to increased convenience: full service became available at any branch of the bank, rather than merely the branch at which the account is kept. But this was bought at a price: anyone's bank details can now be bought from a private eye for a few hundred pounds. It only takes one corrupt teller at each bank, and privacy is lost. Even one careless teller – who wrongly believes that a caller is an insider and divulges personal information over the telephone – is enough. False-pretext phone calls are a large-scale source of information leakage in the UK public sector [1]. One existing government system – the Police National Computer – tries to block them; it has a substantial audit resource in order to check at random whether it is being abused. There is no suggestion in the Framework of a similar arrangement being required elsewhere; yet audit is essential to deter and detect corrupt or negligent practice by staff.

15. The recent proposals to make prison sentences available for those who wrongfully obtain personal information are most unlikely to be a sufficient compensation for the steady removal of internal barriers to lateral information flow. First, the new law would penalise the private eye who obtains information by deception, but not the civil servant who carelessly discloses it. Second, the enforcement record of the Information Commissioner is lamentable. Third, minor privacy offences are not a police priority (especially where the offender and the victim are in different parts of the country). Compartmentation, data separation and constraining systems functionality are crucial to reducing safety, security and privacy risks. Moreover it is only by seriously limiting the functionality and scale of electronic systems that we can have any hope of achieving the high assurance systems that the most sensitive data require.

16. Overall, the Framework will encourage system designers to believe, wrongly, that it is fine to aggregate government information into national databases so long as appropriate authentication takes place at the front end. An example is the electronic health record currently being implemented in the NHS National Programme for IT. Here it is assumed that national-scale records are safe and acceptable, and that patients can be given the means to log on and inspect their own records online, so long as some care is taken with identifying them and issuing them with passwords. This is emphatically not the case. First, it must be assumed that the front-end authentication will be bypassed frequently, and at times on an industrial scale, by attacks involving phishing, pharming, and other variants of deception. Second, the assumption that the front end can take the strain will lead designers to ignore the difficult and subtle issues that were traditionally taken care of in back-end controls. For example, how should a system deal with people who are coerced into logging on and revealing records?

17. Consider a 15-year-old girl who has had a termination of pregnancy and not told her parents. We are aware of two such cases. In one of them, her parents became suspicious and tried to find out via the GP, who managed to prevent them looking at the relevant record; in the other, a relative of the child worked in the health authority and leaked news of the procedure to her family. In scenarios like this, patient privacy – and physical safety – may depend on interposing professional judgment between a request for health record access and the delivery of the record. A blithe assumption that "automation is OK provided there's authentication" may lead to serious harm.

18. Authentication system designers who lack detailed application knowledge will usually be unable to make accurate assessments about back-end controls. The Framework provides a good example at 11.2, which describes a "pseudonymous medical screening service". The writers of this section assume that pseudonyms used (for example) in the management of HIV/AIDS cases are sufficient to protect the identity of patients who present for screening, and therefore that confidentiality concerns are relatively light. However, the "pseudonyms" used for many years in HIV care – postcode plus date of birth plus Soundex code of surname – allow almost all UK residents to be identified (the exceptions being mostly twins living with their parents). The proposed analysis procedure thus returns a result that is completely wrong, as it considers confidentiality compromises of pseudonymous records to be unimportant.

19. The stock response proposed is IL3 client authentication – which will impose costs, but not stop phishing attacks. A better solution might be a code arranged between the doctor and the patient: "Mr Bloggs, we'll text you 'A714' if you're HIV positive and '80FC' if you're negative". But concentrating on the authentication misses the real problem – which is to minimise the suicide risk following a positive diagnosis, by providing counselling and access to suitable support networks. Protocols should be designed around that, rather than being twisted to fit the Framework; and they need space to evolve over time to fit their human environment [2]. Here, following the Framework's recipe just produces a clunky design that diverts attention from the real problem and is likely to put patient safety at risk.

20. In conclusion, this Framework is the sort of approach to security that one might have expected from the public sector ten years ago. It is not fit for purpose in the 21st century. We suggest that the Cabinet Office should, by public tender, invite bids for it to be rewritten.

Professor Ross Anderson
Nicholas Bohm
Dr Brian Gladman
Paul Whitehouse

Foundation for Information Policy Research
14th March 2007

References

[1] RJ Anderson, *Security Engineering*. Wiley, 2001. Available online at `http://www.cl.cam.ac.uk/~rja14/book.html`

[2] MA Blaze, "Toward a Broader View of Security Protocols'." *Security Protocols 2004*, Springer Lecture Notes in Computer Science v 3957 pp 106–132. Available online at `http://repository.upenn.edu/cis_papers/273/`