Healthcare Protection Profile – Comments

Ross J Anderson

Cambridge University
ross.anderson@cl.cam.ac.uk

It is very welcome to see healthcare protection profiles being developed. However, some changes and clarification could make them much more widely acceptable – specifically in Europe, with its rigorous data protection laws.

My background is that I have for a number of years advised various medical bodies on the safety and privacy of health information systems – most recently the British Medical Association and the Icelandic Medical Association; I wrote the security policy of the former [1]. I also organised the first conference on the subject in 1996 [2], and edited a special issue of the Health Informatics Journal in 1998 [3].

1 Scope

The healthcare information which we need to protect includes 'primary' records – the information used in hospitals, clinics or family doctors' offices to support the care of patients directly, and 'secondary' records – billing information, de-identified databases used in research, and so on.

These are two quite different protection tasks. The first is fairly easy and non-controversial, as doctors and patients have more or less shared goals and the problems are amenable to technical solutions involving access control, encryption and similar mechanisms.

The second is much more difficult as it involves genuinely competing interests. Healthcare payers want to retain billing data for ever as it may have marketing or cost control value in the future; patients would prefer it to be deleted as soon as the bills have been paid. Researchers would like access to full longitudonal patient records, if necessary with patient names encrypted; however, linking the healthcare episodes involving an individual patient can make the patient identifiable rather quickly. For this reason, HCFA maintains two sets of de-identified medical records. 'Beneficiary-encrypted' files are assumed to be re-identifiable if anyone makes the effort, and thus available only under fairly strict controls; 'public-access' files are much more thoroughly scrubbed but are correspondingly less valuable to researchers.

The two protection goals should involve different protection profiles; apart from the differing technologies and trust relationships, there are substantial differences in the treatment of secondary medical records between European countries. For example, the UK government collects extensive data centrally to assist in the management of the National Health Service, while attempts to collect almost exactly the same data in Switzerland have been blocked and reversed by their Federal privacy authorities. In addition, the legislative environment for secondary records is different in Europe from America. The EU Data Protection Directive stipulates special protection for personal data pertaining to health (and to matters such as religion, race and sexual preference that can be health related) while US legal protections are looser.

So in my view the draft profile for a 'healthcare provider intranet with limited internet exposure' is a very appropriate first step in developing a suite of healthcare protection profiles.

2 Protection requirements

In what follows I assume that the security target is an integrated medical records system in a hospital.

There has been substantial non-US work on developing security policies for such environments. Policies have been proposed and/or adopted by the British, German and Swedish Medical Associations, the Canadian Health Informatics Association and the EU project SEISMED; unlike with secondary record systems, these policies impose largely similar requirements [4]. A system compliant with the British policy has been implemented and run for several years at three hospitals [5] [6]. So in Europe at least, the protection of patient data in hospitals is a fairly well studied problem with a number of fielded solutions.

The fundamental lesson learned from this experience is that, unlike multilevel secure systems which are designed to prevent information from flowing 'down' (e.g., from Top Secret to Secret to Comfidential), medical systems must be designed to prevent information flowing 'across' (e.g., from ward 1 to ward 2) except where this is authorised.

Thus while the typical rule enforced by a military system might be expressed in lay language as 'a clerk cleared to Secret may read data at Secret or at Confidential, but nothing at Top Secret', so the typical rule enforced by a hospital system is 'a nurse may read the record of any patient who has been on her ward in the previous 90 days'.

There is some variance in how such systems are described. The US military describes them as 'compartmented' (they are used by some agencies to handle intelligence data), a common European term is 'multilateral security'. Rather than arguing over names, we should note that the health systems industry now has considerable experience of building such systems.

Attempts to impose the multilevel security policies familiar from government systems have failed. For example, there was an attempt in the UK to introduce security classification levels into medical systems: the idea was that AIDS databases would be at Secret, normal records at Confidential and administrative records a level corresponding to 'unclassified but sensitive'. This turned out to be unworkable as many administrative records allowed secret facts to be deduced; for example, a record of a patient's visit to an abortion clinic or a clinic for sexually transmitted diseases should clearly be kept secret.

One problem that must have faced the developers of the healthcare protection profiles is that most the available tools and precedents for supporting Common

Criteria activity are oriented towards multilevel secure systems. The authors of the draft profiles appear to have gone along with this flow.

3 Things that need fixing

The main problems with the existing draft profiles, from the European viewpoint, are:

1. There is an assumption that no users are hostile (A.NOEVIL). This is wrong. To a first approximation, all privacy failures in medical record systems result from abuse of authorised access by insiders. This is why the compartmented, or multilateral, approach to security is traditional in medicine: it is an acceptable risk for a family doctor's receptionist to have access to the records of a few thousand patients, but not for such access to extend to all the patients in the healthcare system.
2. It's assumed that the system must defend itself against attempts by unauthorised users to tamper with it (O.SELPRO). This is not enough; it must also defend itself against tampering by authorised users. In other words, it must enforce the security policy independently of user actions.
3. There is still a requirement for multiple levels of data (P.INFOCTRL). This should be removed and replaced with a requirement that access be controlled by role, by ward or by department.
4. The assumption A.RESTRICT simply states that restricted information (such as patient records) should be available 'on a need-to-know basis'. This ducks the question at the heart of any security policy: who defines the need? In the case of healthcare it's even less adequate, as needs don't confer rights. There are many circumstances in which a patient's refusal of consent completely overrides the needs of other parties. For example, a patient diagnosed with HIV in the UK is encouraged to tell his family doctor and all other health professionals involved in his care. A surgeon about to operate on him has a clear need to know the patient's HIV status (so he can put on an extra pair of gloves, etc). Yet if the patient refuses to share his HIV status, this is completely decisive. The additional risk to which the surgeon is exposed is simply an accepted cost of maintaining confidence in the doctor-patient relationship.

   In general, the use of the phrase 'need-to-know' in a security policy sounds an alarm bell to the discerning reader. It often means that the protection requirements have not been thought through, so hand-waving has been resorted to instead. Its appearance in the protection profile will significantly undermine its credibility.
5. Finally, health systems have severe audit requirements: it must be possible at all times to reconstruct the state of the system as it was at any time in the previous six years. This is for both safety and medico-legal reasons. Medical records must persist for at least this period (and for some diseases, such as cancer, for life); and staff accused of malpractice need to be able to reconstruct the data available to them at the time a treatment decision was taken.

4 Conclusion

I welcome the protection profile initiative; it has the potential to consolidate the currently fragmented markets for healthcare informatics systems and products throughout the USA and Europe. However, profiles must take account of established European laws, standards, systems and working practices. The current draft of the profiles is too oriented towards multilevel security. European laws require, and many European systems implement, multilateral security (or compartmented security as it's known in the US government systems community). Only a relatively small number of changes need to be made to the protection profile to accommodate this; but they will greatly enhance its prospects of success on the larger stage.

References

1. 'Security in Clinical Information Systems', BMA (1996); ISBN 0-7279-1048-5
2. 'Personal Medical Information – Security, Engineering and Ethics', Springer (1997) ISBN 3-540-63244-1
3. 'Health Informatics Journal' v 4 no 3/4 (December 1998); ISSN 1460-4582
4. 'Generic security polices for healthcare information systems', S Kokolakis, D Gritzalis, S Katsikas, in [3] pp 184-195
5. 'Implementing access control to protect the confidentiality of patient information in clinical information systems in the acute hospital', I Denley, S Weston-Smith, in [3] pp 174-178
6. 'Privacy in clinical information systems in secondary care', I Denley, S Weston-Smith, BMJ v 318 (15 May 1999) pp 1328-1331