

Comments on the Security Targets for the Icelandic Health Database

Ross Anderson

University of Cambridge Computer Laboratory
Ross.Anderson@cl.cam.ac.uk

Introduction

I have been asked by the Icelandic Medical Association to comment on the following documents written by Admiral Management Services Ltd and made available for public comment by the Icelandic Data Protection Commission:

Security Targets for an Icelandic Health Database by John Arnold (document 7163/T/1);

Approval Process Methodology by Clair Groom (document 7163/T/2).

My comments are as follows.

1 General Comments

The security target document appears to ignore two fundamental problems:

- that the parties against whom the privacy of patients must be protected are the operating license holder, the end users, the system developer, and the sponsor for evaluation;
- that these are for all practical purposes the same organisation (‘Decode’).

This is as far as I am aware a unique situation in the history of secure systems evaluation. There have been other applications in which all the parties are mutually mistrustful; these range from nuclear arms control treaty verification to prepayment electricity meters. However in no application in my experience, or described in the literature, have the threats to one of the stakeholders been so concentrated.

As in systems like treaty verification and metering, third party threats – from outside hackers, disgruntled employees – are almost insignificant compared with the threat of abuse by authorised insiders.

The current custodians of personal health information in Iceland (the doctors and allied professionals) do not trust Decode, and the purpose of having an independent evaluation is to provide assurance that Decode will not abuse its access to, and custody of, clinical and genetic data in a way that discloses personal health information or otherwise contravenes medical ethics.

While I cannot claim to have followed the long argument between the two camps in great detail, I will record that in 1998 I advised the Icelandic Medical Association (IMA) on the proposed database, and pointed out a number of flaws in the proposed protection scheme; and on the 9th December 1999 I met Decode's computer security manager, Úlfar Erlingsson, at his request, at a conference which we both attended in Arizona. This background helps me appreciate the depth of distrust between the IMA and the management of Decode. Although my interest in the proposed database concerns the technical issues, engineering is not performed in a vacuum. It is clear that for the evaluation to meet its stated purpose it should assure doctors that the health database's protection mechanisms will resist attacks by Decode's management.

2 Problems with the Security Target Document

Admiral appears to have made a serious error in the organisational diagram on page 6. Here, the 'end user' is represented as part of the 'database management unit' rather than as part of the 'operating license holder'. The impression given by Dr Erlingsson to the writer was that the DMU would consist of a bank computer vault and the operators would be facilities management staff from the banking industry. Such an arrangement could provide a significant amount of assurance. In New Zealand, which has a similar system, the end users are specially vetted and trained civil servants; their database is held up as an example of good practice. But in the proposed Icelandic system, the end users will work for the license holder, and if they are included in the database management unit then many of the benefits of using neutral third parties will be lost.

The effect of this conceptual error on the security target document is much wider. It suggests that the main protection mechanism is the 'identity removal service' rather than the query layer - where many of the serious problems lie. The query layer is presented as an internal component of the database management unit in this diagram. Later in the text (p 18, 6.5.2.2) it is noted that the data protection commission will have the sole power to approve queries, while elsewhere (at 6.7.4.1 (a)) we find a milder view, namely that the data security committee will set the statistical parameters.

But the query design is probably the most important protection mechanism and involves many deeply technical issues; Erlingsson admitted to me that it may involve developing genuinely novel science. If this is not evaluated, the whole evaluation exercise will be of little effect. The erroneous diagram also gives the impression that separation of duty can be maintained between database management staff and others (7.4.1 (f) assumes that at most one insider will be involved in any attack, but the document avoids developing separation-of-duty issues any further, e.g. at 4.2.4). In short, the document fails to tackle the dual control issues properly.

The second noteworthy fact about the security target is the low levels of assurance it demands - EAL3 for technical mechanisms and EAL1 for the security environment. EAL1 consists of little more than management assurances,

which the IMA is unlikely to accept from Decode. The operating policy and procedures merely need to exist; they can be informal, and they don't have to be reviewed by the evaluator for effectiveness or even consistency (see section 7.2). Yet with the present design, the procedural mechanisms are essentially all that stand between the database and insider abuse.

Evaluating the technical mechanisms to EAL3 is slightly better but not much. It means that vulnerability analysis is taken on trust from the developer, that the evaluator has no access to the source code, that mechanisms are only moderately resistant to attack, and that detected vulnerabilities be merely documented rather than eliminated. With mistrustful principals of whom one is the software developer, an evaluation level of EAL4 (which mandates source code inspection) would be absolutely necessary, and EAL5 (adding a mathematical analysis of the statistical security mechanisms) would be the level required to give practical assurance that the job had been done right.

Admiral's assertion (p 36) that EAL3 is a compromise, appears to support the conclusion which I drew in 1998: that given the proposed application, it was doubtful whether adequate protection mechanisms could be engineered in practice, but in fairness one should allow Decode the opportunity to come up with a design.

Thirdly, the security target document's comments on statistical security don't inspire much confidence that the application, and its already anticipated problems, have been understood.

The summary of attacks (p 9) refers to out of date academic literature (substantially one textbook from 1982). It ignores recent work, including for example Latanya Sweeney's highly relevant data detective research, as well as the concerns raised specifically in the context of the present database about attacks involving genealogical data (such as re-identifying data using kinship patterns).

Three mechanisms are proposed – perturbation, query set size control and subsetting – which are infeasible in the proposed application. Indeed, the writer suggested them in 1998 to Decode's technical director and discussed them again with Dr Erlingsson in December 1999. The Decode view as expressed to me accepts that ranging – e.g. describing a patient as 'female aged 70-74 years' rather than 'female born 21/11/1929' may be useful. Query set size control may have some limited uses. But the data are largely diagnoses and kinship relationships rather than numbers, so it's not clear how perturbation can be used; and Decode's desire to use as large a population as possible rules out subsetting as a general mechanism.

But subsetting is the only control contemplated by Admiral as a defence against tracker attacks, which are the most widely known and documented threat to statistical databases (9.2.4 (j) (i)). So the system will be undefended against this type of attack. In short, the security target is predicated on primary defence mechanisms which are already agreed to be unworkable.

Next, audit. The target requires a vast amount of audit material to be produced (including all operations on protected objects – 6.2.2.3 (g)), but is largely silent on how this material is to be used. Is it to be scrutinised by the

license holder (in which case it's valueless for preventing attacks by license holder management) or by the data protection commission, in which case what tools have to be constructed to analyse it? How will such tools be evaluated? At 6.2.5.2 we read 'the TSF shall provide the audit records in a manner suitable for the user to interpret' – this surely cannot be right as it's the user's actions which are the target of the audit activity. And how will all this be related to the threats and the security objectives?

Further defects in the audit section include 6.2.2.3 which says (inter alia) that imminent security breaches must be audited, while 6.2.6 says that the TSF shall not record any information in the audit records that could lead to disclosure of personally identifiable data. This is inconsistent. 6.2.6 would preclude logging unsuccessful logon attempts: these very often contain the password in the user-name field. (There are more complex problems such as the creation of covert channels out of the database, but this simple example should do to persuade the reader that the information flow control issues haven't been thought through).

At 6.2.11 there's a more active attack: access administrators can use the system without leaving an audit record if the audit trail is full. Introducing an error that will fill up the audit disk is easy, and this is the sort of attack that a knowledgeable insider is likely to carry out. For this reason, operating systems such as MVS and the Berkeley distributions of Unix take steps to insulate their audit trails against abuse by administrators.

Next, crypto. One of the curious things about the original design presented to the IMA by Decode in 1998 was the use of encryption to protect the database, when clearly inference controls were what was actually needed. It became clear that Decode did not at that time have anyone who understood inference controls (a deficiency since addressed by hiring Dr Erlingsson). Yet the same old ineffective crypto design appears in the security target document.

For example, Decode proposed in 1998 that personal identifiers be protected by one-way encryption or by public key encryption. But with a population of well under a million, an opponent with access to the encryption key can simply encrypt the names of the entire population and look for a match. To provide privacy protection, the system would have to use random padding – but this would render the encrypted identifiers useless for their intended purpose. This was pointed out to Decode in 1998, and is one of the criticisms expressed in a publicly available English language document which the IMA commissioned and which is available on the web ("The DeCODE Proposal for an Icelandic Health Database", Ross Anderson, Oct 20, 1998; <http://www.cl.cam.ac.uk/#Med>).

Other issues concerning encrypted data in statistical databases don't seem to be understood. Data that has had patient names encrypted ('beneficiary-encrypted' data, in American jargon) is still personal health information as it can usually be re-identified easily from the context. This is the whole point of using inference controls rather than (or as well as) encrypting names. Yet at 6.4.1.4 we read that encrypted data is to be treated as system data and is only covered by the FPT* protection mechanisms (start-up self-test and transaction atomicity). This is a gross error.

Next, there are some details.

1. 6.6.3.1 requires that passwords be at least 16 characters long with upper and lower case alpha, and numeric, values. This ensures that operators must write them down, and so is of little value without rules about how they are to be stored (e.g., in a safe in an office that's locked when not occupied);
2. 7.4.4 misses out FAU* and FDP* whose strength of function is much more important than the crypto. If an insider can defeat the operating system access controls using a stack smashing script downloaded from a site such as rootshell.com, then he could assume administrator privileges and perform arbitrary actions without the audit trail recording them. That is far more serious than a hypothetical keysearch attack in 70 years' time;
3. 8.1.5 (f) refers to a data warehouse containing unencrypted health records. I was not briefed on this during my visit to Iceland. Have the IMA been briefed on it?
4. 9.2.4 (j) (iii) claims that insertion attacks aren't possible. This is untrue. New patient records are added all the time, and many of them may be completely known to the license holder (for example, those patients who've already consented to share fully identified data).

3 Conclusion

The final issue is whether the evaluator, and indeed the consultant contracting for the security target, is likely to acquire and maintain the trust of both Decode and the Icelandic Medical Association. The need for this is succinctly stated in section 2.4 of the 'Approval Process Methodology' document.

I am concerned that the security target presented by Admiral has so many serious errors that it completely fails to inspire confidence, and indeed calls into question whether Admiral possesses the skills required to evaluate the security of a statistical database. Admiral seems not to understand the operating environment or the system structure; they haven't read even the publicly available English language literature on the system; they don't seem familiar with the relevant science; the primary protection mechanism which they propose – sub-setting – is agreed by both parties to be largely irrelevant; elementary errors about the use of cryptographic primitives, which were pointed out over a year ago, have been repeated; and patient-encrypted data are considered to need only cursory protection. There are also a number of technical errors such as failing to give access controls the consideration which they require.

In my view, an evaluation proceeding according to this security target would very unlikely to inspire confidence and would thus be of little value. The Data Protection Commission should appoint a new consultant with the appropriate skills and start afresh.