An Update on the BMA Security Policy

Ross Anderson

Cambridge University Computer Laboratory
Pembroke Street, Cambridge CB2 3QG
ross.anderson@cl.cam.ac.uk

Abstract. In this article, we attempt to step back from the current dispute between the BMA and the government and describe it as a whole. We give a brief account of the origins and development of the BMA security policy and guidelines. We then summarise the feedback so far, and discuss its practical implications (which were the focus of official objections). Experience of pilot projects and systems overseas shows that many of the problems can be solved fairly easily by available technology.

The policy has clarified things significantly, and we now see that the remaining 'hard' problems are unavoidably political. They pit long established patient rights and professional privileges against the NHS's Information Management and Technology Strategy, which directs healthcare computing investment away from clinical systems to build a series of databases that will make personal health information available centrally to administrators. Our investigation of this has been slowed (though not thwarted) by systematic official obstruction, which suggests that administrators are uncomfortably aware of the ethical problems.

1 Introduction

In late 1994 and early 1995, the British Medical Association (BMA) repeatedly asked officials of the UK National Health Service (NHS) about encryption of data on a new data network that was being planned. The assurances received were less than convincing. They included the claim that there was no encryption expertise in Britain, and the even more bizarre claim that encryption could not be introduced until the network was in place, as the network itself would be needed to distribute the keys [65] [66] (it was later learned that encryption proposals had been spiked at the request of the intelligence community). I was therefore contacted and asked to speak to the BMA's Information Technology Committee (as it now is) on the 8th March.

On looking at the documents that the government had supplied to the BMA on security in the proposed network [50] [51] [52] [53] [54], it was clear that something was wrong. The government assumed that the main additional threat from connecting clinical computer systems together would come from outside 'hackers' — a view common enough in the popular press but not held by people with experience of the field.

The likelihood that data will be abused depends on its value and on the number of people who have access. Connecting systems together increases both these risk factors at the same time. An example is given by personal financial information, which in many countries is no longer private: as any bank teller can access any account at most banks, an illegal data broker needs only a small number of sources to cover most of the population's finances [44] [64]. The prospect of medical records suffering a similar fate is alarming, and the controls proposed by the government would have been unable to prevent this.

The NHS argument was that for 'security' reasons, all clinical data would have to be carried on their private network that was being set up by a contractor, BT. Organisations wishing to connect to it (and all significant healthcare providers would be forced to) would have to sign a 'Code of Connection' promising not to connect their systems to any other network [54]. But however convenient the Code for BT's business at a time of rapidly growing competition and falling costs for data network services, it would provide no protection against the majority of attackers who would, we believed, come from inside the system rather than from outside.

Our concerns were first communicated to the government in detail in a letter from the BMA on the 21st March 1995. This questioned the assumptions that the NHS network could be kept separate from the Internet and that encryption was infeasible; it also pointed out inconsistencies in the NHS security policy. It received a testy response. Thus, on the 31st May, the BMA Council supported a resolution from the IT working party that the problems with the threat model, security policy and architecture would "need to be addressed as a matter of urgency by the NHS Executive or use of the NHS Wide Network would be boycotted for the transmission of identifiable patient data by doctors concerned about confidentiality".

So we prepared a detailed critique [4] of the NHS threat model, security policy and architecture and presented it to senior officials on the 8th June 1995. At that time, we fully accepted the bona fides of the NHS Executive and aimed to help them revise their security policy and architecture documents to be acceptable. In the world of security, it is common practice that one party advances a design and another tries to find holes in it. Such third party evaluation is a standard industry practice, and is mandatory in many government systems in Britain, the EU [39] and elsewhere.

2 The Gathering Storm

We were not to know it at the time, but the NHS Executive had projects underway to build systems that are in serious conflict with medical ethics as understood by both doctors [31] [32] and patients [17] [36] [59]. If security rules are adopted that enforce this traditional view, then these systems will require significant changes (which we discuss below).

So, with the benefit of hindsight, it is not at all surprising that the response we received from the NHS Executive was limited to nitpicking [47], ad hominem attacks, diversionary tactics (such as the recent report on encryption [77]) and delay.

This surpassed the script of "Yes Minister". For example, at a meeting called on the 26th June to present their response to our critique, officials claimed that we would have to wait for the NHS to settle its confidentiality policy — a document that had been stalled for some 15 years, and the most recent version of which (in August 1994 [14]) had been roundly rejected by clinical professions, patients and the Data Protection Registrar. So the Association went public with its concerns; these were summarised in an article that appeared in July [3].

By then it had become rather clear that the government was determined on a tactical rather than constructive response. Our intelligence sources reported a determination to implement the Code of Connection and deal with objectors by obfuscation, delay and diversion; the strategy was to field the network and present it as a fait accompli. Typical of the tactics used in this period was a letter in September that sought to query the minutes of the 26th June meeting and wished a further meeting in November to discuss them [48]. Also in September, a senior IMG official claimed at a conference that our criticisms had been completely misguided, as the primary purpose of the NHS network was to provide leased lines between hospitals that would cut phone bills!

In spite of these Fabian tactics, the foundations of the government's position were removed one by one. The erroneous initial assumption — that the main additional threat from networking would come from outsiders — was repudiated in a report commissioned by NHS managers from the government's own expert body, the CCTA [55]; the four level 'classification' of data that formed the intellectual core of their security policy and justified their architecture was next to go [67]; yet officials stuck adamantly to their 'Code of Connection'. In vain we pointed out the practical problems that would arise — Addenbrookes' Hospital, for example, shares its network infrastructure with Cambridge University. These objections were ignored.

More senior officials became involved, and their tactics became steadily more reckless. A very senior medical officer wrote in August that the government would press ahead with its Code of Connection and hoped that the BMA objections could be dealt with later [75]; when we objected to the use of the network for clinical information, he claimed that Item-of-Service claims were not personal health information and that contract minimum data sets were 'of course coded' [76]. For the benefit of readers not familiar with NHS systems, a typical Item-of-Service claim is for the supply by a general practitioner of contraceptive care, and that a typical contract minimum data set is for an episode of hospital treatment. I was personally lost for words that one of the government's most eminent doctors could hold unworthy of protection the identities of under-age girls taking the pill or obtaining pregnancy terminations in NHS hospitals.

On the 8th December 1995, the Code of Connection was issued, despite senior officials having given assurances to the BMA on the same day that this would not happen [27]; it was promptly denounced by the Association [28]. The Code, together with supporting documents such as the IS Security Reference Manual [56], continued to use the security assumptions and arguments that had already been discredited by the government's own experts.

We pointed this out and on the 13th December a senior official wrote to the Association:

> You have included references to IMG project documents. These are project working papers provided to project members ... you will see that they are classified "Restricted: Management" ... please therefore delete the references [67].

No assurances of confidentiality had been sought by the government, or given by the Association, when these documents were originally supplied.

3 The Policy is Commissioned

By September 1995, the BMA had become convinced that the NHS Executive either would not or could not draw up an acceptable security policy, and so on the 7th October the BMA Council asked me to do this. My goal was not to rewrite the traditional ethics of the profession, but to translate them into a concise set of rules that would provide a clear and unambiguous basis of communication between patients, clinicians and policymakers on the one hand, and computer system builders on the other.

There already existed two well understood security policy models to provide some inspiration. The first is the Bell-LaPadula policy, used by the world's armed forces, under which an official cleared to 'secret' should be able to see documents classified 'secret' and below, but nothing at 'top secret' or above. In other words, information only flows upwards, and never downwards, through a hierarchy of security levels [9]. The second is the Clark-Wilson policy that was developed to formalise good practice in banking and bookkeeping systems, and which lays down a number of rules to enforce controls such as dual control and audit [19]. But neither of these would do for clinical information, the basic principle of which is expressed by the General Medical Council [31] as:

> Patients have a right to expect that you will not pass on any personal information which you learn in the course of your professional duties, unless they agree.

Thus our goal is patient control of data access, rather than an access hierarchy that reflects an organisational command structure. It is privacy, that empowers the patient, rather than confidentiality, that empowers the organisation. This distinction is already familiar to medical ethicists: in English law, the privacy of medical records is founded on the rights of the patient while the confidentiality of social work records is based on the rights of the local authority that employs the social worker [24]. However, it was less familiar in the computer security world, as previous security models (including both Bell-LaPadula and Clark-Wilson) had been driven by organisational rather than privacy concerns.

So how could privacy — the principle of patient control — be encapsulated in a compact set of rules that would be easily understood by patients and clinicians, but sufficiently precise for system builders?

The BMA also commissioned guidelines. The idea was that the policy would be normative — it would state where we should be in a few years' time — while the guidelines would tell the working doctor how to protect her patients (and herself) from the immediate threats. One might think of the policy as the long-term treatment plan, and the guidelines as a bandage to stop the bleeding.

Developing the policy was a fascinating experience. The main primary sources used to elucidate the GMC position were the books by Somerville on medical ethics [72], and by Darley, Griew, McLoughlin and Williams on clinical confidentiality [24]. These provided the background material on what problems arise in practice, and how the clinical professions expect them to be dealt with. The pioneering study of electronic patient records by Griew and Currell [30] was also useful; it showed how complex it is to build a policy model for a record containing components to which different combinations of clinicians would have access, and motivated the search for a simpler framework.

The key idea was to assume that each record would have a unique access policy. That is, we would treat a lifetime's medical history as an accumulation of records, each of which was completely accessible to a the same set of users. Thus the general record might be available to everyone in a practice or care team, while a note on a treatment for depression might be open only to the doctor who treated it (and to the patient). This greatly simplifies things, and has the virtue of reflecting actual clinical practice.

By early November 1995, a first draft of the policy was circulated, and was significantly refined by a number of discussions. Among the most helpful were presentations to the BMA's IT and Ethics committees; we also shared the early drafts with software suppliers so that any practical objections could be raised, and with the NHS Executive, whose contribution at the time was negligible. These meetings took place during November and December 1995.

The final versions of the policy and guidelines were written over the New Year holiday and shipped in early January 1996 [5] [6]. The core of the policy is contained in nine principles, which are appended. A period of public consultation ensued, of which this workshop is the logical culmination.

4 Post Publication Feedback

The feedback on the security policy, from both institutions and individuals, has been roughly of three kinds. Firstly, the majority of responses have been strongly supportive (e.g., [26]). A common comment has been that the work brings clarity to a subject that many had for years found to be confusing, and that while its principles may not all be achievable at once (or even at all in some legacy systems), it shows where we should be going. At least one medical school has dicussed incorporating the policy into its curriculum.

The second kind of response has come from officialdom and its sympathisers, who emphasise 'practical' objections to the policy. This became amusingly clear at a meeting with officials on the 6th February at which a senior official claimed that the principles would be impractical, as the notification requirements would be too onerous. We informed him that we would be resolving this question by conducting a trial at a number of general practices. He then said that although the principles might work in general practice, they might be impractical in a hospital setting. A clinician present asked whether he was suggesting a trial in the context of GP-hospital links and he replied that that would be an adequate trial. We promptly agreed and minuted the agreement. In a later letter, he complained that this was still not wide enough to test the principles' practicality [68].

The rest of the criticisms — the interesting and useful kind — are made up of a large number of observations by various parties, but with a number of recurring themes.

1. A number of clinicians have argued that integrated hospital systems can bring important safety benefits; they might help prevent the tragedies that can happen when records go astray (as many paper records do [1] [2]). The point is also made that at some hospitals, as many as 70% of admissions are accident and emergency, so there is little scope for compartmentation between clinical departments [63]. When one asks advocates of integrated hospital systems how to control the aggregation threat that arises when many hospital staff can see data on many patients, and which will become much worse if hospitals are connected together into a network, the suggestions include:
   - forego NHS networking as insufficiently important;
   - allow only a small number of trusted staff to copy records from one hospital to another, and audit them closely;
   - remove general access to records of patients who are not currently receiving treatment. A typical acute hospital might have files on a million people, but only a few percent might be active (as in- or out-patients) at any one time. Only a small number of trusted library and admissions staff would have the ability to restore a record to 'active' status;
   - our suggestion was to use a technology such as active badges [73] to track hospital staff, and prevent (or investigate) accesses to the records of patients in other departments or wards. It turns out that a similar system is used in some US hospitals but based on departmental groups of terminals. Staff who access another department's records face questioning and possible disciplinary action [23];
   - educate the public to change their expectations of medical privacy.

   In any case, the practicality of securing hospital information systems is an open question, with some contributors foreseeing serious problems [63] and others not [35]. Resolving this will be an empirical matter, and may involve some exceptions to the policy — an issue which we will discuss further below.
2. One of the most trenchant criticisms came from a senior member of the computer security community, Gus Simmons (who was for many years the senior scientist at Sandia National Laboratories, whose responsibilities include the security of the US nuclear arsenal). He argued that it is not adequate to secure an electronic system to the same level as the paper system it replaces, as critical social controls are removed.

With a paper records system, an attacker can always grab a file from someone else's office, but this activity is counter to social taboos, and is fraught with risk that the occupant might return unexpectedly. But when records are placed on a computer, anyone who can get access through his terminal will not appear to a passer-by be doing anything wrong. Thus he may feel that he is committing at most a very minor misdemeanour. So electronic record keeping systems should have very strong auditing and intrusion detection systems; a deterrent that must be publicised and credible [71].

3. As an intrusion detection mechanism, Simmons suggested that whenever anyone looked at a patient's record but did not bill the patient for her time, then it should be investigated as a prima facie abuse. This would harmonise the patient's interest in privacy and the hospital management's interest in maximising its revenue.
4. Similar ideas were suggested independently by Ulrich Kohl [43]. His development is somewhat more general and shows that context-based access controls can be implemented with with quite general parameters.
5. On the other extreme, the policy has been criticised for not emphasising that computerised medical records have the capability to be much more secure than paper records [63]. We have never disputed this as a possibility — but have still to see a really secure electronic medical record system fielded.
6. A number of contributors worried about the extent to which access control lists would have to be micromanaged, and whether this would turn out to be a serious burden given the large number of record fragments that can pertain to one individual [62]. In fact, given the signal-to-noise problems, might it not turn out to be unfeasible?

   Our view was that the great majority of individuals can be dealt with using a default access control list, containing a group such as 'all GPs working in the practice', and that only a small number of highly sensitive records would require exceptional treatment with an access control list containing only the treating doctor and the patient himself. Nonetheless this was felt to be an extremely important question, and in consequence was one of the points investigated in a trial of the principles carried out in a number of general practices. Some early results are described in the paper by Alan Hassey and Mike Wells [36].
7. A number of contributors objected to the restrictions on aggregating patient data. A typical comment was "There is no doubt that general aggregated data, such as immunisation uptake, has been beneficial to the common good ... system linkage or networking has, I would suggest, been poorly planned and perhaps somewhat hurried ... however I do feel that it is inevitable and that the benefits will ultimately outweight the perceived pitfalls" [60]. Several groups opined that with de-identified data it might be extremely difficult to obtain information such as analysis of readmissions to hospitals [22] [70]. This is a very important point, and one on which US contributors also had much to say; we will deal with it fully in a later section.
8. Tom Rindfleisch made the point, with which we fully agree, that informed

consent should not be sought at the stressful point of critical need, but in advance, like a living will [62]. A related point is made by the German information security agency: that for consent to be meaningful, systems must be designed so that people who refuse to use part or all of them, or to grant some information access, do not lose their right to care as a result [10]. The German case referred to a health smartcard; it is unclear what would happen if someone needing hospital treatment in the UK refused permission for their personal health information to be entered on the Clearing system, and discussions with officials have elicited only the vague suggestion that perhaps the hospital would simply foot the bill for treatment itself "as a one-off".

9. Some members of the computer security community objected to principle 9 (the Trusted Computing Base), on the grounds that it is a part of the security engineer's basic intellectual environment. However, the BMA policy talks to clinicians as well as technicians, so we feel it is appropriate. No matter how the document is written there will be parts that some section of the audience feels to be superfluous.
10. Some writers preferred 'fuzzier' statements of the security policy goals and want it to be more 'patient centred' [61]. We remain unmoved. A security policy is like a scalpel: it must be clean and sharp rather than warm and furry. As for the buzzword 'patient centred', systems so described often seem to be a cover for transferring the primary record from the GP to a health authority, a hospital or an insurance company. We are satisfied to have upheld the principle of patient control.
11. A number of computer companies complained that the security functionality required was so different from that offered by their current products that expensive redevelopment would be necessary [11]; while the Association of the British Pharmaceutical Industry asked for some of the principles to be made less 'draconian' [74].
12. We received quite a lot of input on practical solutions used elsewhere, e.g. German cancer registries [12], the New Zealand registry system [58], and similar registries registries implemented in Denmark [45] and proposed in Norway [13]. The point was also made — from experience with HIV programmes in the USA — that apart from neonates, the date of birth is clinically irrelevant and should be suppressed in clinical systems, thereby reducing the likelihood of harmful linkages being constructed with other systems at some later date [15].

We noted above that some exceptions to the policy may have to be made, e.g. for accident and emergency staff. This does not of course invalidate the policy. Even policies such as Bell LaPadula and Clark-Wilson fail to cover their application areas completely. In a bank, for example, there are typically about twenty roles which cannot realistically be subjected to dual control, such as the chief executive, the chief systems programmer, the computer security manager and the chief dealer. Such people simply have to be trusted, despite the fact that the trust occasionally turns out to be misplaced.

This is well understood in the security comunity. The security policy sets a yardstick; system builders get as close to it as they economically can; the shortcomings are examined during the evaluation process; and so when the system is presented by the contractor to the customer, he can make an informed decision on whether to accept the residual risk or send the system back for redevelopment. The policy does not eliminate residual risk, but rather quantifies it and enables a prudent judgment to be made about it.

That kind of benefit should materialise once people start using the policy to build systems. Meantime the main benefit is clarity. The policy has enabled us to work through the logical consequences of the GMC's ethical principle — that patients should have control over access — in much greater detail than ever before, and apply it as a test to many fielded and proposed systems.

Previously, discussions had tended to set a rather poorly defined 'patient confidentiality' against an equally poorly defined 'public interest' that was often described vaguely in terms of research benefits but was all too often a front for attempts to increase official power and control. However the policy, and its followup in the GP pilot, brought us to identify the tension between privacy and safety as the best way to express the trade-offs from the patient's point of view.

Leaflets distributed as part of the GP pilot reflect this, and the GP pilot also enabled us to identify the flows of information from general practice to health authorities, for such purposes as item-of-service claims and cervical screening, as one of the few problems with implementing the policy and guidelines in general practice [42] [36].

It also brought to our attention that there are potentially major problems with de-identification of the data in statistical databases. However, before we explore this, it is appropriate to mention the feedback received from conferences in the USA.

5 A Lesson from America

A version of the BMA security policy was accepted for presentation at the IEEE Symposium on Security and Privacy at Oakland, which is the premier conference on computer security, and we submitted a condensed version that incorporated much of the early feedback [8]. After the paper was presented (on the 7th May), there was a panel discussion at which an academic, a doctor and a representative of the healthcare computing industry presented their views of the policy. Then, on the 10th May, the policy was presented again at a workshop in Washington at which doctors, lawyers, rights activists and congressional staffers discussed the issues from a US viewpoint.

The main lesson learned from this trip was that the real privacy problem in the USA comes from the claims databases operated by the insurance companies that pay for most US healthcare. These databases are coming to replace the casenotes in the doctor's office as the primary record for many Americans; the convenience of having a lifetime's record in one place outweighs the fact that these records were not generally designed for clinical use.

The sidelines the security debate. US hospital computer systems have much greater variety than their UK counterparts, and their level degree of security also varies widely. But there is a feeling that, since patient records can be obtained by almost anyone from the insurance industry, why should more money be invested in making hospital systems any better?

One of the Oakland speakers revealed that his company sees the seven million records kept by its health systems division as a major business asset, and would strongly resist any attempt by legislators or others to restrict the ways in which this could be used to produce revenue. As we noted in the policy, this business structure has led to practices that would be considered highly abusive in the UK. For example, forty percent of insurers disclose personal health information to lenders, employers or marketers without customer permission [18]; over half of America's largest 500 companies admitted using health records in personnel decisions [16]; and US firms are regularly taken over for the value of the medical records under their control. Indeed, most Americans are coming to feel that these practices are worrying, and a quarter have personal experience of abuse [33].

This has led to a number of bills being introduced or proposed at both state and federal level, and is the subject of papers elsewhere in this volume. Here we will remark that aggregated records make a tempting target. For example, at the Washington meeting a district attorney discussed his use of medical records in criminal investigations. He saw nothing remiss in issuing a subpoena for insurance company files that he thought might be helpful — and insurance files (being considered financial rather than health records) enjoyed no special privilege.

Another serious aspect of claims-based longitudinal records is that they are not accurate. It is common to 'inflate' diagnoses so as to be able to claim higher fees, so that, for example, non-specific chest pain will be recorded as ischaemic heart disease. This might be qualified as a tentative diagnosis in the clinical notes, but as the 'unified computer record' supersedes this, the false diagnosis may prevail. It was mentioned that some 20% of alleged clinical facts in the computer record were wrong; if this is even the right order of magnitude, then the risks to health are significant.

The social effects of insurance-driven data aggregation are also becoming understood. At the Washington meeting, a primary care physician told us that over the last twenty years, US patients have moved from complete trust in their family doctor to a much more guarded relationship, in which patients suppress facts that are potentially embarrassing or harmful. The risks of this should also be clear.

6 Could it Happen Here?

The standard response of NHS officials on being told of information abuses in the United States is 'it couldn't happen here'. Yet the US trip focussed our attention on the threat from the construction of large databases of personal health information. There had already been signs that all was not well.

An internal presentation by the NHS Executive to the effect that there should be a unified electronic patient record, shared by everyone in the NHS, had already caused concern — to the extent that we had confronted senior officials on the 31st January and asked whether the real goal of the IM&T strategy was to construct a series of centralised databases, each covering a different aspect of health care but which would together contain essentially all personal health information on every NHS patient — in effect, nationalising the country's medical records using contract data as the Trojan Horse for the project.

This was stoutly denied. Officials categorically assured us that the abstracts of the contract data that were kept centrally were not only de-identified, but also unlinkable — separate episodes concerning the same patient could not be correlated. This was claimed to be a property of the HES data formats. We accepted these assurances and asked for a copy of the HES data specifications; we were promised a copy (which never turned up). Incidentally, the claim that central databases contain only episode data is still being repeated by senior officials [49].

The next stimulus came in February 1996 from an HIV data collection project. This was presented as an attempt to improve planning for HIV sufferers, who at present can self-refer to any hospital in the UK rather than having to go through their GP. As a result, officials suspected that the 18,000 registered sufferers represented only about 12,000 actual patients, and wanted to know if budgets could be cut. A form was sent out to all GPs and genitourinary clinics demanding details of all patients receiving treatment [46]. In addition to clinical information, this demanded that the patient be identified by date of birth, post-code and the 'Soundex' code of their surname[1]; the instructions for generating a Soundex code have the curious final line 'Note: it is very helpful if you can give the initial of the first name as well'.

This information was being chased up, and handled, by employees of district health authorities, rather than being sent directly to the Public Health Laboratory Service. The development of regional databases is also mentioned in the protocol, but without detail. When these concerns were made public, a consultant epidemiologist at the laboratory claimed that "Somebody who does not know what the Soundex code is would have no possibility of guessing the identity" [37] — hardly reassuring given that the Soundex system is public and that the patient's name and data of birth are present on the form!

Meanwhile, it was pointed out that HIV status was already encoded in the contract minimum data set, as were codes for other sexually transmitted diseases, abortions and fertility treatment [34].

The next stimulus was in March 1996 when a study of the NHS Executive's IM&T strategy commissioned by the BMA's IT Committee reported that

> The changes to the flows and management of health information will, when completed, represent the most fundamental and challenging changes to the practice of medicine ever [57].

[1] essentially, this means the initial letter and the next three consonants

7 Linkable After All!

Given this background, the US experience caused us to stop and reexamine the overall pattern, which entailed looking at the ultimate repository and beneficiaries of the large quantities of information that the Information Management and Technology Strategy sets out to gather. We had still not received the HES data definitions that the government had promised in January, so these were now obtained otherwise.

This led to the shocking discovery that the categorical assurances which we had received about the HES data were completely false. The records in this database contain the full postcode, date of birth and sex [25]. So with a few exceptions (such as twins living together, and students in colleges) the patients are easily identifiable and the episodes are linkable. In fact, it is unclear what their value would have been otherwise, as one of their avowed functions is to assess hospital readmission rates.

This contributed to an impression that the Department of Health has for some time worked to create a set of central databases with details of every episode of care in the country. If this is the case, then no doubt knowing that it would be controversial, they have tried to do it by stealth.

This impression is not dispelled by ministerial assurances. An MP had set down the following parliamentary question about the Clearing service, the central system for settling health care payments between purchasers and providers, and which also skims off the HES data for central government [21]:

> To ask the Secretary of State for Health, ... on what basis (contractor's) employees or managers will have access to personal data?

The government replied [38]:

> Their managers and employees are contractually bound to maintain the confidentiality of data passing through the Clearing Service, and will have no access to it.

This is intrinsically implausible to a computer security person (surely the system administrators will have access?), and when we obtained a copy of the Clearing system documentation we found that according to its security policy, staff with 'a direct operational functional requirement' would have access to personal health information, while access to information that had been 'de-identified' (i.e., with the name and address removed but with the postcode and date of birth still presumably present "shall be available to all Users for healthcare business purposes, subject to receipt by the Contractor in writing of rules imposed by the Data Protection Registry".

So it appears that our initial fears were well founded. In addition to the Clearing and HES systems mentioned above, there are databases in existence for prescriptions and planned for community care and data collected from general practice. Meanwhile, the government states that matching of official data will be allowed by officials investigating welfare fraud. Is it reasonable to hope that access will be denied to police, customs, tax officials, and indeed every official who can plead a 'need to know'?

8 The Way Forward

Even under the charitable explanation — that the government's actions are the result of blunder rather than malice — we face the unpleasant fact that the databases that are to support research and business information have been made identifiable by using, as the primary key, the combination of date of birth and postcode.

Quite apart from the privacy issue, this will cause both safety and reliability problems. Firstly, although most of the population can be uniquely identified in this way, a minority cannot — twins living at the same address, for example, and students in halls of residence (for whom the capture-recapture problem in probability theory ensures that if over 23 people of the same age are living at the same address, then at least two of them are likely to share the same date of birth). Thus, if in the absence of a paper record, an accident and emergency team digs out a HES record and acts on the information it contains, then there is a small but significant probability that they will be using the wrong person's data.

Another problem is that the linkage of records will be broken when patients move. This will distort hospital readmission statistics, as it can be assumed that changes of address will be correlated with illness (assuming illness to be correlated with unemployment, divorce and homelessness).

We would therefore recommend that, as a matter of urgency, the National Health Service — together with all its information systems contractors — cease and desist from using (date of birth, postcode) as a primary database key.

Instead, the techniques developed in Denmark and Germany should be used. Each healthcare provider submitting data centrally should use a pseudonym, whose linkage to the patient is unknown to outsiders. For example, one might pass the name and date of birth through a hash function such as SHA1 [69], together with a key unique to the provider, and take as many bits of the result as necessary to fill the fields in question. If the use of techniques that smack of cryptography is to be forbidden, then one can simply generate the pseudonyms at random (and take care to protect the file that links them to patient identities).

Either way, the use of systematic pseudonyms would lessen the risk of the wrong record being used, and also reduce the loss of information linkage — many address changes are local, and these patients remain with the same provider even when their postcode changes. It would also bring these systems into line with the established RCGP/GMSC guidance:

> no patient should be identifiable, other than to the general practitioner, from any data sent to an external organisation without the informed consent of the patient [40]

Such simple measures will not completely solve the problem, as people with access to the databases might infer a patient's identity from knowledge of part of their clinical history — as we pointed out in the policy. However it would eliminate the most serious problem and build a foundation on which further inference controls could be constructed (see, e.g., [29]).

It will also not tackle the problem that once large central databases exist, then there will be pressure for researchers to use these for reasons of economy. Official control of these databases then might have a negative effect on paradigm-breaking research. How readily would the establishment grant access to future scientists making unconventional claims, such as a link between Helicobacter Pylori and ulcers, or between Chlamydia and coronary heart disease?

9 Conclusions

The Secretary of State for Health is reacting to the success of the BMA's campaign against the NHS wide network by focussing health IT spending on precisely the objectionable components of the IM&T strategy (the NHS wide network, the new NHS number, the NHS wide Clearing service) at the expense of clinical systems [20]. This is strange for a conservative minister presumably alert to the dangers of centralisation and aware that the market in health systems is perfectly capably of matching willing buyers with willing sellers without the need for a central civil service department to set up national monopolies in service sectors which already have competitive provision.

We have advanced a possible explanation for the urgency. The government is building a series of linkable databases — Clearing, HES, PPA, registers for HIV, diabetes and other expensive diseases, and future databases covering primary and community care. These will eventually aggregate under central control all personal health information of significance. Although they are represented as being 'anonymised', they are nothing of the kind. The project may be justified internally as 'creating an electronic patient record shared throughout the NHS', but externally the picture is different. Officials are so sensitive about it that they have systematically obfuscated and delayed; it has taken over a year for us to dig down through successive layers to the heart of the problem.

But it is not necessary for these databases to contain identifiable information. In fact, as we have shown, replacing the current primary database key of postcode and date of birth with a one-way hash function of name and date of birth would bring tangible safety and accuracy gains.

If the database building project proceeds without controls of this kind, it can only be construed as a political attempt to centralise personal health information for state purposes. If that comes to pass, we may expect that health privacy in Britain will go the way of America. The papers in this volume by observers of the American scene give us some idea what to expect then.

Appendix — the BMA Security Policy Principles

Principle 1: Access control Each identifiable clinical record shall be marked with an access control list naming the people or groups of people who may read it and append data to it. The system shall prevent anyone not on the access control list from accessing the record in any way

Principle 2: Record opening A clinician may open a record with herself and the patient on the access control list. Where a patient has been referred, she may open a record with herself, the patient and the referring clinician(s) on the access control list

Principle 3: Control One of the clinicians on the access control list must be marked as being responsible. Only she may alter the access control list, and she may only add other health care professionals to it

Principle 4: Consent and notification The responsible clinician must notify the patient of the names on his record's access control list when it is opened, of all subsequent additions, and whenever responsibility is transferred. His consent must also be obtained, except in emergency or in the case of statutory exemptions

Principle 5: Persistence No-one shall have the ability to delete clinical information until the appropriate time period has expired

Principle 6: Attribution All accesses to clinical records shall be marked on the record with the subject's name, as well as the date and time. An audit trail must also be kept of all deletions

Principle 7: Information flow Information derived from record A may be appended to record B if and only if B's access control list is contained in A's

Principle 8: Aggregation control There shall be effective measures to prevent the aggregation of personal health information. In particular, patients must receive special notification if any person whom it is proposed to add to their access control list already has access to personal health information on a large number of people

Principle 9: Trusted Computing Base Computer systems that handle personal health information shall have a subsystem that enforces the above principles in an effective way. Its effectiveness shall be subject to evaluation by independent experts.

References

1. *'Setting the Records Straight — A Study of Hospital Medical Records'*, Audit Commission,, June 1995
2. *'For Your Information — A Study of Information Management and Systems in the Acute Hospital'*, Audit Commission,, July 1995
3. "NHS wide networking and patient confidentiality", RJ Anderson, in *British Medical Journal* v 310 no 6996 (1 July 1996) pp 5–6
4. *'NHS Network Security'*, RJ Anderson, 30th May 1995
5. *'Security in Clinical Information Systems'*, RJ Anderson, published by the British Medical Association, January 1996; also available from `http://www.cl.cam.ac.uk/users/rja14/#Med`
6. "Clinical system security: interim guidelines", RJ Anderson, in *British Medical Journal* v 312 no 7023 (13 Jan 1996) pp 109–111

7. "Patient Confidentiality — At Risk from NHS Wide Networking", RJ Anderson, *to appear in Proceedings of Healthcare 96, March 96*
8. "A Security Policy Model for Clinical Information Systems", in *Proceedings of the 1996 IEEE Symposium on Security and Privacy* pp 30–43
9. DE Bell, LJ LaPadula, *'Secure Computer Systems: Mathematical Foundations'*, Mitre Corporation report ESD-TR-73-278
10. *'Chipkarten im Gesundheitswesen'*, Bundesamt für Sicherheit in der Informationstechnik, Bundesanzeiger 4 May 1995
11. *Submission from HBO & Company*, J Baker
12. B Blobel, *this volume*
13. *'Pseudonymous Medical Registries'*, E Boe, Norwegian Official Report 1993:22
14. *'Draft guidance for the NHS on the confidentiality, use and disclosure of personal health information'*, N Boyd, DoH, 10 August 1994
15. V Brannigan, *personal communication*
16. "Is your health history anyone's business?" McCall's Magazine 4/95 p 54, reported by M Bruce on Usenet newsgroup comp.society.privacy, 22 Mar 1995
17. "Confidentiality of medical records: the patient's perspective", D Carman, N Britten, British Journal of General Practice v 45 (September 95) pp 485–488
18. "Who's reading your medical records?" *Consumer Reports*, Oct 94 pp 628–632
19. "A Comparison of Commercial and Military Computer Security Policies", D Clark, D Wilson, in *Proceedings of the 1987 IEEE Symposium on Security and Privacy* pp 184–194
20. "Dorrell urges refocus over NHS technology", in *Computer Weekly* (30/5/96)
21. Parliamentary question, H Cohen, 3/4/96
22. *'Security in Clinical Information Systems'*, submission from J Crown, President, Faculty of Public Health Medicine, to BMA, 29/2/96
23. R Cushman, *this volume*
24. *'How to Keep a Clinical Confidence'*, B Darley, A Griew, K McLoughlin, J Williams, HMSO 1994
25. *NHS Data Manual, Technical Modules Volume 1 and 2*, 1996
26. *Submission from the Society of Occupational Medicine*, D Dean, 12/4/96
27. "New Guidance on Computer Security Issued", *DoH press release*, 8/12/96
28. "BMA warns doctors about government guidance on computer security", *BMA press release*, 11/12/96
29. *'Cryptography and Data Security'*, DER Denning, Addison-Wesley 1982
30. *'A Strategy for Security of the Electronic Patient Record'*, A Griew, R Currell, IHI, University of Wales, Aberystwyth, 14/3/95
31. *'Good Medical Practice'*, General Medical Council
32. *'Confidentiality'*, General Medical Council
33. "Privacy and Security of Personal Information in a New Health Care System", LO Gostin, J Turek-Brezina, M Powers et al., in *Journal of the American Medical Association* v 20 (24/11/93) pp 2487–2493
34. "Contract minimum dataset includes confidential data", in *British Medical Journal* v 312 (20/1/96) p 185
35. *(HISS presentation to BMA IT Committee, 24/4/96)*
36. A Hassey, M Wells, *this volume*
37. "HIV code prompts debate on privacy", P Hagan, in *Hospital Doctor* (29/2/96) pp 16
38. Parliamentary reply, J Horam, 16/4/96

39. *'Information Technology Security Evaluation Criteria'*, EU document COM(90) 314 (6/91)
40. "GMSC and RCGP guidelines for the extraction and use of data from general practitioner computer systems by organisations external to the practice", Appendix III in *'Committee on Standards of Data Extraction from General Practice Guidelines'* Joint Computer Group of the GMSC and the RCGP, 1988
41. "Nurse Jailed for Hacking into Computerised Prescription System", in *British Journal of Healthcare Computing and Information Management* v 1 (94) p 7
42. S Jenkins, *this volume*
43. U Kohl, *this volume*
44. "Your Secrets for Sale", N Luck, J Burns, *The Daily Express*, 16/2/94 pp 32–33
45. *Private conversation with Peter Landrock*
46. " 'Soundex' codes of surnames provide confidentiality and accuracy in a national HIV database", JY Mortimer, JA Salathiel, Communicable Disease Report v 5 no 12 (10 Nov 1995) pp R183–R186
47. Senior IMG official, letter to BMA, 22/6/95
48. Senior IMG official, letter to BMA, 7/9/95
49. Senior IMG official, talk on Radio Northampton, 11.10, 12/6/96
50. *'Information Systems Security: Top level policy for the NHS'*, IMG document 2009 (b)
51. *'NWN Threats and Vulnerabilities'*, 5 April 1995, IMG document NWNS/T1.22
52. *'NHS-wide networking: data security policy'*, IMG document NWNS/T3.3
53. *'NHS wide networking security architecture'*, 3 April 1995, IMG document NWNS/T1.21
54. *Security Guide for IM&T Specialists'*, 3 April 1995, IMG document NWNS/T5.11
55. *'NHS/CCTA Internet Security Report'* version 1.3
56. *'NHS IS Reference Manual'*, December 1995
57. *'A Members' Guide to the Intended Goals and Purposes of the IM&T Strategy'* R Neame, 3/3/96
58. R Neame, *this volume*
59. "GP Practice computer security survey", RA Pitchford, S Kay, Journal of Informatics in Primary Care (September 95) pp 6–12
60. *letter from DR Price to BMA, 28/5/96*
61. M Rigby, *this volume*
62. *Presentation to IEEE Symposium on Security and Privacy 96*, T Rindfleisch, 7/5/96
63. R Roberts et al, *this volume*
64. "For Sale: your secret medical records for £150", L Rogers, D Leppard, *Sunday Times* 26/11/95 pp 1–2
65. Senior NHS Executive official, letter to BMA, 20/12/94
66. Senior NHS Executive official, letter to BMA, 15/2/95
67. Senior NHS Executive official, letter to BMA, 13/12/96
68. Senior NHS Executive official, letter to BMA, 12/2/96
69. *'Applied Cryptography'*, B Schneier, second edition, Wiley 1995
70. *Response on behalf of Conference Information Group*, Prof. M Severs
71. GJ Simmons, *personal communication, 1996*
72. *'Medical Ethics Today — Its Practice and Philosophy'*, A Sommerville, BMA 1993
73. "The Active Badge Location System", R Want, A Hopper, V Falcao, J Gibbons, in *ACM Transactions on Information Systems* v 10 no 1 (January 1992) pp 91–102
74. *Submission on behalf of the ABPI*, F Wells

75. Senior NHS medical officer, letter to BMA, 15/8/95
76. Senior NHS medical officer, letter to BMA, 17/11/95
77. *'The use of encryption and related services with the NHSnet'*, prepared by Zergo Ltd for NHS Executive; document NHSE IMG E5254

This article was processed using the LaTeX macro package with LLNCS style