

Information Governance in NHS's NPfIT: A Case for Policy Specification

Moritz Y. Becker

Microsoft Research

Cambridge, CB3 0FB, United Kingdom

moritzb@microsoft.com

Keywords: Electronic Health Record, Patient Confidentiality, Trust Management, Authorisation Policy, Access Control

Abstract

Purpose The NHS's National Programme for IT (NPfIT) in the UK with its proposed nation-wide online health record service poses serious technical challenges, especially with regard to access control and patient confidentiality. The complexity of the confidentiality requirements and their constantly evolving nature (due to changes in law, guidelines and ethical consensus) make traditional technologies such as role-based access control unsuitable. Furthermore, a more formal approach is also needed for debating about and communicating on information governance, as natural-language descriptions of security policies are inherently ambiguous and incomplete. Our main goal is to convince the reader of the strong benefits of employing formal policy specification in nation-wide electronic health record (EHR) projects.

Approach Many difficulties could be alleviated by specifying the requirements in a formal authorisation policy language such as Cassandra. The language is unambiguous, declarative and machine-enforceable, and is based on distributed constrained Datalog. Cassandra is interpreted within a distributed Trust Management environment, where digital credentials are used for establishing mutual trust between strangers.

Results To demonstrate how policy specification can be applied to NPfIT, we translate a fragment of natural-language NHS specification into formal Cassandra rules. In particular, we present policy rules pertaining to the management of Clinician Sealed Envelopes, the mechanism by which clinical patient data can be concealed in the nation-wide EHR service. Our case study exposes ambiguities and incompletenesses in the informal NHS documents.

Conclusions We strongly recommend the use of trust management and policy specification technology for the implementation of nation-wide EHR infrastructures. Formal policies can be used for automatically enforcing confidentiality requirements, but also for specification and communication purposes. Formalising the requirements also reveals ambiguities and missing details in the currently used informal specification documents.

1 Introduction

The UK National Health Service's (NHS) National Programme for Information Technology (NPfIT) is one of the most ambitious and challenging ongoing IT projects world-wide: at the heart of the project lies the development of a nation-wide online service, also known as the Spine, enabling health professionals and patients to access electronic health records (EHR) from anywhere, anytime. Its success is of utmost importance, not only in face of its huge potential benefits to UK public healthcare, but also because of the risks involved. It is the world's most expensive IT project in the public sector [1], so failure would result in huge financial losses. But more importantly, as the project involves aggregating vast amounts of patient-identifiable confidential data in a widely-distributed networking environment, failure to protect the data appropriately may even lead to the ruin of people's lives [2, 3, 4].

The security and confidentiality requirements (henceforth referred to as *information governance*, in accordance with NPfIT terminology), as set out in the Output-Based Specification (OBS) [5] and a subsequent NPfIT document [6], are highly challenging and beyond the capabilities of current access control technologies including role-based access control (RBAC). The challenges of information governance are manifold:

- The proposed Spine will be extremely large, holding life-long EHRs of 50 million patients.
- It is widely-distributed: central services such as the Spine or the Personal Demographic Service (PDS) will interact with tens of thousands of local clinical information systems, all of which must comply to both local and global policies.
- The rules governing access to patient-identifiable information are complex and reflect the trade-off between patient confidentiality, usability, and legislative constraints. Neither traditional discretionary nor mandatory access control nor RBAC are sufficiently expressive to capture complex policies such as Legitimate Relationships, Sealed Envelopes or patient consent management.
- The requirements are mandated by laws, official guidelines and ethical positions that are prone to change. Such changes have to be implemented quickly and consistently across all local and central systems and services.
- Many issues are complex, novel, not well understood, and/or controversial. Communicating such issues in plain English alone (be it within NPfIT consultation groups or between NPfIT and their IT suppliers) is problematic: any such description is inherently incomplete and ambiguous.

Many of these difficulties could be alleviated by adopting the trust management approach, introduced by Blaze et al. in 1996 [7, 8]. Here, access authorisation is based on digitally signed credentials containing principals' attributes rather than mere name bindings, thereby enabling mutual strangers to share resources in large distributed systems. Furthermore, authorisation rules are specified in a high-level policy language that is machine-enforceable. A number of such languages have been proposed in the past few years, including KeyNote [9], SDSI/SPKI [10], Ponder [11], and Lithium [12]. Some languages such as QCM [13], SD3 [14], Binder [15], RT [16] and Oasis [17] also support automated retrieval of credentials over the network. Cassandra [18, 19] additionally features automated trust negotiation [20, 21], a process in which sets of credentials are exchanged to build up mutual trust.

It is exactly applications as complex and challenging as the NPfIT Spine that trust management and policy specification technology aims at. In this paper, we purport to convince the reader of the strong benefits (or even necessity) and feasibility of employing policy specification in the development of any nation-wide EHR service. For this purpose, we will demonstrate how information governance rules for NPfIT can be specified in Cassandra, a particularly flexible and expressive authorisation policy language. To keep this paper short, we will focus on rules relating to "Clinician Sealed Envelopes" that allow clinicians to conceal data from patients.

The benefits of policy specification are laid out in §2. An overview of the policy rules pertaining to Sealed Envelopes is given in §3. We conclude with a discussion of ambiguities in the specifications, lessons learnt and future work.

2 Why Policy Specification?

The official OBS document [5] specifies that RBAC [22] should be used for authorisation. In RBAC, users are assigned to roles, and roles are associated with privileges that can be exercised by activating a role. Role hierarchies express seniority relations and facilitate privilege inheritance. Roles thus provide an extra level of indirection that make access control administration of enterprises more

scalable. However, RBAC alone is not sufficiently expressive for information governance on the Spine. The reason is that the processes of role membership and privilege assignment, and role activation and deactivation are subject to a multitude of constraints, or *rules*. Some of these rules are well-known idioms such as separation-of-duties constraints: for example, a user must not log on as a clinician and a patient at the same time. But there are many such rules, and they come in different variants, so it is not sufficient to extend RBAC by a fixed number of built-in constraints. It would be even worse to hard-wire such rules into the implementation in an ad-hoc way. Instead, we propose the use of a policy language to explicitly specify such rules.

Policy specification should be used in the NPfIT project for two purposes. Firstly, it should be used as a *communication aid* for issues on information governance, within NPfIT as well as between NPfIT and their suppliers, and for specification purposes. Informal descriptions should be supplemented by formal and precise policy rules. NPfIT would benefit from this approach for the following reasons:

- A formal policy is unambiguous, precise and yet concise.
- It is much more concrete and specific than current, natural-language specifications such as the OBS, but abstracts away irrelevant implementation details.
- Policy rules are ideal for presenting alternatives and making clear the (often subtle) differences.
- As formal policies are machine-enforceable, it is easy to build a simulator application for stakeholders (NHS experts, suppliers, clinicians and patient representatives) to “play around with” and thus to explore the consequences of the policy or policy alternatives empirically.
- A formal policy precisely specifies the compliance criteria for suppliers. Furthermore, a policy engine fed with the policy could produce randomised test cases for regression testing of implementations.
- Policy rules are amenable to formal analysis, with which security properties of the policy can be proven mathematically.

Secondly, policy specification and *automated policy enforcement* should be used to govern access control in actual systems. The policy and the policy engine would act as a protective layer between the user interface and the restricted system functions and data. Central services such as the Spine or the PDS as well as local systems would benefit from this approach for the following reasons:

- The approach is consistent with established software engineering and security engineering principles: access control should be independent of system implementation.
- It is by far more maintainable and cost-effective: changes in information governance requirements can be implemented quickly and easily simply by amending the high-level policy. There is no need to change and recompile any source code, and to shut down and restart the service.
- Checking and certifying compliance of systems would become simpler as only the high-level policy has to be checked.

3 Policy for Sealed Envelopes

We have written Cassandra policy rules for the Spine and related services that encompass the whole range of requirements of the OBS Version 2 [5] concerning access to patient-identifiable records, including job role certification, workgroup management, Legitimate Relationships, clinicians and patients concealing data, explicit consent and patients’ agents [23]. Because of the inherent ambiguity and incompleteness of the OBS, we had to fill in many details and interpret the specification using common sense. Moreover, the policy comprises 375 rules. Surely, some of our rules will not be in accordance with the actually intended meaning. Furthermore, some points will have changed since the OBS was published in August 2003.

We will now present a revised fragment of our policy, namely those rules pertaining to Clinician Sealed Envelopes, the mechanism by which clinicians can conceal EHR data. The revision is largely based on [6] from January 2005, a draft document that concretises some of the points left unclear in the OBS and in some details also departs from it. The presented fragment has a more manageable size and suffices to demonstrate how policy specification can be applied to NPfIT. The fragment omits rules concerning Legitimate Relationship and consent management, patient agents, third-party consent, RAs and the PDS. The rules also differ from those in the complete, but unrevised policy in [24].

Cassandra is based on Datalog with Constraints known from Constraint Logic Programming (CLP) and is thus relatively intuitive. The following rules are accompanied by textual explanations, but some basic familiarity with Cassandra may be helpful. For lack of space, we have to refer the reader to [18, 19] for an introduction to Cassandra. In the following, the section numbers starting with §730 refer to the OBS section on Information Governance.

3.1 Clinician Log-on

Access to system functions and data is role-based (§730.9), and job roles such as “Clinician” are associated with more specific specialties or areas of work such as “GP” or “Psychiatry”. Before a user may access any system function, she has to log on with one (and only one) job role (§730.12.10). Professional users will authenticate themselves using a smart card. It is conceivable that the smart card will contain digital job role credentials signed by NHS-approved registration authorities (RA) (§730.21, also cf. sections on authentication, registration and RBAC in [6]). Personal users (patients) will access the services through the NHS HealthSpace web portal.

For the purpose of illustrating the Sealed Envelope policy fragment, we will only describe the rules relating to logging on with the job role “Clinician”, corresponding to the parameterised Cassandra role `Clinician(org, area)`. As expected, the parameter `org` refers to the organisation the clinician is working in, and `area` to her area of work. Of course, if a clinician works in several organisations or in several different areas, she will be able to activate the `Clinician` role with different parameters (though not simultaneously: §730.12.10).

The rule C01 uses the special¹ `canActivate` predicate to state that a user `cli` may activate the `Clinician` role with parameters `org` and `area` (thus logging on with the corresponding job role) if

- the Spine has a copy of a credential (either in its policy or because `cli` has submitted it along with the activation request) issued by some `ra` (this is specified by the prefix `ra` in front of the predicate), and certifying that `cli` is an NHS clinician working in the organisation `org` and in the work area `area`;
- `ra` is an NHS-approved registration authority for the organisation `org` (cf. rule RA01 below);
- `cli` has not activated any other job role or personal role (the defining rules for the user-defined predicate `no-main-role-active` are not presented here) (§730.12.10);
- and finally, the credential is still valid, i.e., the current time lies between the credential’s start and end dates (§730.24.7).

(C01)
`canActivate(cli, Clinician(org, area)) ←`
`ra.is-certified-NHS-clinician-cert(cli, org, area, start, end),`
`is-registration-authority(ra, org),`
`no-main-role-active(cli),`
`Current-time() ∈ [start, end]`

Alternatively, if the RA-issued clinician credential is not available to the Spine (because the clinician has neither submitted it nor is it stored in the Spine’s policy), the Spine will request an RA-issued credential from the clinician’s policy (this is specified by the prefix `cli▷ra` in front of the predicate) (C02):

(C02)
`canActivate(cli, Clinician(org, area)) ←`
`cli▷ra.is-certified-NHS-clinician-cert(cli, org, area, start, end),`
`is-registration-authority(ra, org),`
`no-main-role-active(cli),`
`Current-time() ∈ [start, end]`

The user-defined `is-registration-authority(ra, org)` predicate used above is defined by Rule RA01. It is satisfied by a credential signed by the NHS and asserting that `ra` is an RA for the organisation `org`. Furthermore, the credential must be still valid.

(RA01)
`is-registration-authority(ra, org) ←`
`NHS.is-NHS-registration-authority(ra, org, start, end),`
`Current-time() ∈ [start, end]`

A user who has logged on with a `Clinician` role can also log out by deactivating her own role activation (C03), specified using the special predicate `canDeactivate`:

(C03)
`canDeactivate(cli, cli, Clinician(org, area)) ←`

¹Apart from user-defined predicates, Cassandra has six special predicates, namely `canActivate`, `hasActivated`, `canDeactivate`, `isDeactivated`, `permits` and `canReqCred` that are used by the access control engine to handle role activation and deactivation, and action and credential requests.

3.2 Clinician Sealed Envelope

Clinicians may seal off data from the patient in exceptional circumstances (see §730.49, and the section on Clinician Sealed Envelope in [6]). Both the OBS and [6] note that the “granularity” of patient data that may be sealed off is subject to further national guidance. We circumvent this uncertainty by assuming that a patient’s record is divided into “items”, without further specifying the granularity of items.

A clinician cli can conceal an item $item$ pertaining to a patient pat ’s record by activating the role $\text{Concealed-by-clinician}(who, pat, item, start, end)$ (note the use of a role that represents a state rather than a simple job role). The parameters $start$ and end represent a validity period for the Sealed Envelope (§730.51.8). The parameter who is a four-tuple $(cli, org, area, group)$ indicating that cli works for organisation org in work area $area$ and is member of the workgroup $group$ that has a Legitimate Relationship with pat . This information is needed later in the deactivation rules CC05 and CC06. Rule CC01 requires cli to have a matching Legitimate Relationship with the patient, enforced by the user-defined predicate $\text{is-treating-clinician}(pat, cli, org, area, group)$. We omit the rules governing Legitimate Relationships. The rule also requires the clinician to have activated the **Clinician** role. This is in accordance with §730.12.10 (system functions can only be accessed when exactly one job role has been chosen).

```
(CC01)
canActivate(cli, Concealed-by-clinician(who, pat, item, start, end)) ←
    hasActivated(cli, Clinician(org, area)),
    is-treating-clinician(pat, cli, org, area, group),
    who = (cli, org, area, group)
```

How do we test whether a patient pat ’s record item $item$ has been concealed by a clinician? The user-defined aggregation predicate $\text{count-concealed-by-clinician}(n, pat, item)$, defined by Rule CC02, can be used to count the number n of matching **Concealed-by-clinician** roles for this item. Therefore, the rule governing patient access to their own records (not shown here) can make use of this predicate to enforce the constraint $n = 0$, i.e. the constraint that no clinician has concealed the item.

```
(CC02)
count-concealed-by-clinician(count(x), pat, item) ←
    hasActivated(x, Concealed-by-clinician(who, pat, item, start, end)),
    Current-time() ∈ [start, end]
```

The following rules govern the conditions for deactivating **Concealed-by-clinician** roles, i.e., for removing Clinician Sealed Envelopes. Any clinician who has activated that role herself can also deactivate it (CC03):

```
(CC03)
canDeactivate(cli, cli, Concealed-by-clinician(who, pat, item, start, end)) ←
    hasActivated(cli, Clinician(org, area))
```

The patient’s GP cli can also remove Sealed Envelopes even if they have been imposed by another clinician $cli2$ (CC04). A GP is a clinician with work area parameter set to the constant **GP**.

```
(CC04)
canDeactivate(cli, cli2, Concealed-by-clinician(who, pat, item, start, end)) ←
    hasActivated(cli, Clinician(org, GP)),
    is-treating-clinician(pat, cli, org, GP, group)
```

Furthermore, a Caldicott Guardian (information security officer) of an organisation org can remove Sealed Envelopes imposed by clinicians of that same organisation (CC05). The projection function π_2^4 selects the second element from a four-tuple (recall that the second element of the four-tuple who is the organisation).

```
(CC05)
canDeactivate(CG, cli2, Concealed-by-clinician(who, pat, item, start, end)) ←
    hasActivated(CG, Caldicott-guardian(org)),
    π_2^4(who) = org
```

The rules CC03, CC04 and CC05 are not based on the OBS as it does not comprehensively specify Sealed Envelope removal. CC06 allows the removal of a Sealed Envelope by clinicians working in the

same workgroup as its creator (§730.51.10):

```
(CC06)
canDeactivate(cli1, cli2, Concealed-by-clinician(who, pat, item, start, end)) ←
  hasActivated(cli1, Clinician(org1, area1)),
  is-treating-clinician(pat, cli1, org1, area1, group1),
  who = (cli2, org2, area2, group2),
  group1 = group2
```

The following rule (CC07) is included for ensuring consistency and makes use of the special predicate `isDeactivated`: if the patient's registration is cancelled (e.g. if the patient has deceased), all relevant Clinician Sealed Envelopes are removed automatically:

```
(CC07)
isDeactivated(x, Concealed-by-clinician(who, pat, item, start, end)) ←
  isDeactivated(y, Register-patient(pat))
```

4 Discussion and Conclusion

Informal policy descriptions as in the OBS or [6] are necessarily ambiguous and incomplete. Formal policy specification exposes such ambiguities and also provides a technique for expressing and comparing alternative interpretations.

In specifying the Spine policy, it was particularly unclear to us how much autonomy will be given to personal users through HealthSpace. There is consensus to give patients read access to their record as well as the right to annotate items. But it is unclear whether patients can, for example, directly request items to be put into a Sealed Envelope. Similarly, it is not clearly specified whether they will be able to directly give consent to data sharing, appoint agents acting on their behalves, and manage Legitimate Relationships with clinicians. It is also unclear whether data relating to a third party can be accessed by the patient if the third party “unlocks” it by registering consent.

Another under-specified issue concerns the withdrawal of privileges. It is common sense that users who log on with a particular job role can also log off again (e.g. Rule C03). However, it is not specified whether there are other users or conditions that can cause a job role to be immediately deactivated. For example, if the clinician certification issued by an RA is prematurely revoked, is it required that the clinician is logged off immediately (if she is currently logged on), or is it sufficient to bar her from logging on to the system the next time? Who exactly can remove a Sealed Envelope, and under which conditions? As many of our rules for removing Sealed Envelopes (CC03–CC07) had to be invented, they may not capture the behaviour that is implicitly intended by NPfIT architects, but could be a starting point for carving out a more detailed specification. Permission withdrawal is also relevant in the context of consent and Legitimate Relationships. Can consent (one-off consent, authenticated express consent, third-party consent) be withdrawn, and if yes, by whom? Under which circumstances can Legitimate Relationships be terminated?

Many ambiguities are subtle. For example, can clinicians seal off items that they cannot view themselves? It is easy to see that the presented rules allow clinicians to seal off such items, although the intended (but unspecified) rule may well forbid that. This alternative behaviour can be implemented by replacing Rule CC01 by the following rule:

```
(CC01')
canActivate(cli, Concealed-by-clinician(who, pat, item, start, end)) ←
  hasActivated(cli, Clinician(org, area)),
  is-treating-clinician(pat, cli, org, area, group),
  permits(cli, Read-item(pat, item)),
  who = (cli, org, area, group)
```

Subtle but important details like this one are easily overlooked if not using a formal specification notation. Using formal specification in the actual implementation to automatically enforce the policy can also make the system more secure: a straightforward implementation approach may be to enforce access restrictions in the user interface directly, e.g. by displaying the seal-off option only for those items the clinician can view. Such a design is bad not just because of its obvious maintainability problems, but also because it is insecure if the user interface can be bypassed and the Spine accessed directly.

Of course, informal descriptions are not without merits. They serve as a useful starting point in the design and refinement process, and any formal policy specification should be accompanied by natural-language descriptions, just as the source code for a piece of software should contain comments.

In our attempt to formalise the NPfIT Information Governance policy, we have found use cases and scenarios (e.g. [25, 6]) to be particularly useful.

In this paper, we have argued that NPfIT Information Governance requirements are too complex for traditional access control technologies including RBAC. We would like to urge EHR designers and implementers to consider policy specification as a tool for expressing and communicating design issues precisely, and also to enforce Information Governance in a real implementation in a scalable and maintainable way. We have demonstrated the advantages and feasibility of the policy approach by expressing the requirements on Clinician Sealed Envelopes in Cassandra, a formal policy language.

Acknowledgements I would like to thank Nick Gaunt for extremely helpful discussions, and Malcolm Oswald for providing me with valuable NPfIT documents. Much of this work was done during my time in the Computer Laboratory at the University of Cambridge. I thank Peter Sewell for helpful discussions, comments and support.

References

- [1] Sean Brennan. *The NHS IT Project: The biggest computer programme in the world... ever!* Radcliffe Publishing, 2005.
- [2] Ross Anderson. Patient confidentiality — at risk from NHS-wide networking. In B. Richards and H. de Glanville, editors, *Current perspectives in healthcare computing*, pages 687–692. Weybridge: BJHC Books, 1996.
- [3] Ross Anderson. Information technology in medical practice: Safety and privacy lessons from the United Kingdom. *Medical Journal of Australia*, 170:181–185, 1999.
- [4] Anthony Browne. Lives ruined as NHS leaks patients' notes (25/06/00). *The Observer*, 2000.
- [5] National Health Service, UK. Integrated Care Records Service: Output based specification version 2. 2003.
- [6] Malcolm Oswald. Information governance in the NCRS – material for NHS presentations – DRAFT v.05. *National Programme for Information Technology*, January 2005.
- [7] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *IEEE Symposium on Security and Privacy*, pages 164–173, 1996.
- [8] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. The role of trust management in distributed systems security. In *Secure Internet Programming*, pages 185–210, 1999.
- [9] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. The KeyNote trust management system, RFC 2704.
- [10] Carl M. Ellison. SPKI requirements, RFC 2692.
- [11] Nicodemos Damianou, Naranker Dulay, Emil Lupu, and Morris Sloman. The Ponder policy specification language. In *Policy Workshop*, 2001.
- [12] Joseph Y. Halpern and Vicky Weissman. Using first-order logic to reason about policies. In *CSFW*, pages 187–201, 2003.
- [13] Carl Gunter and Trevor Jim. Policy-directed certificate retrieval. *Software - Practice and Experience*, 30(15):1609–1640, 2000.
- [14] Trevor Jim. SD3: A trust management system with certified evaluation. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 106–115, 2001.
- [15] John DeTreville. Binder, a logic-based security language. In *IEEE Symposium on Security and Privacy*, pages 105–113, 2002.
- [16] Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130, 2002.
- [17] Jean Bacon, Ken Moody, and Walt Yao. A model of OASIS role-based access control and its support for active security. *ACM Transactions on Information and System Security*, 5(4):492–540, 2002.
- [18] Moritz Y. Becker and Peter Sewell. Cassandra: distributed access control policies with tunable expressiveness. In *IEEE 5th International Workshop on Policies for Distributed Systems and Networks*, pages 159–168, 2004.
- [19] Moritz Y. Becker and Peter Sewell. Cassandra: Flexible trust management, applied to electronic health records. In *IEEE Computer Security Foundations Workshop*, pages 139–154, 2004.

- [20] William H. Winsborough, Kent E. Seamons, and Vicky E. Jones. Automated trust negotiation. In *DARPA Information Survivability Conference and Exposition*, volume 1, pages 88–102, 2000.
- [21] William H. Winsborough and Ninghui Li. Towards practical automated trust negotiation. In *Proceedings of the 3rd International Workshop on Policies for Distributed Systems and Networks*, pages 92–103, 2002.
- [22] Ravi Sandhu. Rationale for the RBAC96 family of access control models. In *Proceedings of the 1st ACM Workshop on Role-Based Access Control*, 1997.
- [23] Moritz Y. Becker. Cassandra: Flexible trust management and its application to electronic health records. Technical Report UCAM-CL-TR-648, University of Cambridge, Computer Laboratory, 2005.
- [24] Moritz Y. Becker. A formal security policy for an NHS electronic health record service. Technical Report UCAM-CL-TR-628, University of Cambridge, Computer Laboratory, April 2005.
- [25] Nick Gaunt. Confidentiality and consent: Use cases applicable to shared electronic health record. *S&W Devon ERDIP Project*, 2003.