

A Distributed Requirements Management Framework for Legal Compliance and Accountability

Travis D. Breaux and Annie I. Antón
North Carolina State University
Raleigh, North Carolina, USA
{tdbreaux, aianton}@ncsu.edu

Eugene H. Spafford
Purdue University, CERIAS
West Lafayette, Indiana, USA
spaf@purdue.edu

Abstract

Increasingly, new regulations are governing organizations and their information systems. Individuals responsible for ensuring legal compliance and accountability currently lack sufficient guidance and support to manage their legal obligations within relevant information systems. While software controls provide assurances that business processes adhere to specific requirements, such as those derived from government regulations, there is little support to manage these requirements and their relationships to various policies and regulations. We propose a requirements management framework that enables executives, business managers, software developers and auditors to distribute legal obligations across business units and/or personnel with different roles and technical capabilities. This framework improves accountability by integrating traceability throughout the policy and requirements lifecycle. We illustrate the framework within the context of a concrete healthcare scenario in which obligations incurred from the Health Insurance Portability and Accountability Act (HIPAA) are delegated and refined into software requirements. Additionally, we show how auditing mechanisms can be integrated into the framework and how auditors can certify that specific chains of delegation and refinement decisions comply with government regulations.

1. Introduction

National and international standards, regulations and laws impose restrictions on business practices to achieve societal goals, such as improving corporate accountability in financial markets or ensuring the privacy of medical records in the healthcare industry. Mature standards and regulations describe specific personnel responsibilities that cut across several business units and require comprehensive documentation to demonstrate how personnel decisions implement standards and regulations. Furthermore, certification boards and government auditors impose penalties on organizations to motivate corrective action in the event of non-compliance. While organizations are often held accountable, recent legislation such as

Sarbanes-Oxley¹ (SOX) and the Health Insurance Portability and Accountability Act² (HIPAA) in the U.S. shift liability towards personnel, imposing fines and prison sentences on individuals for their actions that contribute to non-compliance.

Government regulations usually require a set of artifacts that demonstrate and account for personnel actions taken to comply with the law – the matter of *accountability*. Because business practices often include a significant human factor (e.g., people implementing policies in a potentially ad-hoc fashion), compliance with standards and regulations is complicated by pressures on human performance (e.g., increasing profits, decreasing costs). In large organizations, software systems that support these processes can provide increased compliance assurance by supporting software controls that restrict what actions personnel can perform with oversight under the law. In effect, software provides a means to enable business practices while limiting the improper use of resources that would otherwise violate the law.

Organizations are in need of mechanisms to help assure that operational practices comply with standards and regulations. According to two Ernst and Young surveys of nearly 1,200 international organizations, compliance with regulations and policy surpassed worms and viruses as the primary driver of information security policy from 2004-2007 [21, 22]. Moreover, companies are often required by regulatory rules, such as HIPAA [25], to implement policies and procedures that comply with the law.

To meet the needs of operational controls, we propose a distributed requirements management framework. Our framework provides a transparent and accountable method of ensuring that obligations in standards and regulations are implemented by functional software requirements and software controls. In our framework, personnel satisfy obligations by refining them into functional requirements or by creating new obligations that are delegated to others. Through delegation and

¹ U.S. Public Law 107-204, 116 Stat. (2002)

² U.S. Public Law 104-191, 110 Stat. (1996)

refinement, personnel will contextualize their obligations incurred from standards and regulations using their own business knowledge and goals. Recording the personnel decisions to delegate and refine obligations improves accountability, because each decision can be evaluated and compared against best practices. Auditors can certify these decision chains to demonstrate that, at least at a specific point in time, organizations complied with the intent of the law. Furthermore, as policies change, organizations can re-evaluate their decisions to delegate and refine their legal obligations in a framework that dynamically dispatches these changes to personnel responsible for accommodating these changes. While this work has not yet been validated in practice, it has been mathematically validated using Allow [26] and we believe the framework is relevant, effective, and feasible.

In our previous work, we analyzed privacy policies in healthcare and finance [9, 10], HIPAA regulations [11, 13] and organizational security policies [12] to identify policy elements required to align systems with policies and regulations. In each of these studies, we identified a need to manage obligations in a single, distributed framework. Our proposed framework builds upon this need by managing obligations through delegation and refinement with special focus on the needs of auditors.

This paper is organized as follows: in Section 2, we consider a simple scenario to motivate our framework; in Section 3, we present the framework formalism and definitions; in Section 4, we instantiate our framework by elaborating on an application in the healthcare domain; in Section 5 we discuss related work; in Section 6 we discuss requirements for a tool supporting our framework with our conclusion in Section 7.

2. Requirements Scenario

Consider a scenario in which a Chief Security Officer (CSO) has been assigned the high-level security goal (a non-functional requirement) NFR_1 = “to ensure that corporate information is secure.” The CSO implements NFR_1 by assigning several new non-functional requirements including NFR_2 “ensure computer-based communications are confidential” to his IT security manager in charge of network security. The IT security manager responds by identifying all modes of “computer-based communications” relevant to satisfying her new obligation. As a result, the manager identifies internal web, instant-messaging and e-mail servers among others that use TCP/IP network connections to share information among internal systems. The manager, with both authority over who administers these servers and knowledge of available security mechanisms in these systems, implements her obligation by assigning new functional requirements

including FR_1 “ensure web servers use SSL for internal connections” and FR_2 “ensure mail servers use TLS for internal connections” to relevant system administrators across different departments. A system administrator responsible for administering a mail server running Linux receives FR_2 and implements the requirement with a series of configuration directives that he applies to the system: FR_1 = “install latest OpenSSL libraries,” FR_2 = “compile and configure Sendmail with TLS support”, FR_3 = “generate X.509 certificates for Sendmail,” etc.

At each level in the delegation hierarchy, a manager knows *what* goal his staff member must achieve but the manager may not have the technical knowledge to know *how* his staff will achieve this goal. Each obligation is owned by someone who is ultimately accountable for that obligation and the decisions to refine an obligation are also recorded. Tracing permissions and obligations through ownership, delegation and refinement allows managers and auditors to quickly and effectively identify how and why vulnerabilities are addressed to reduce risk of non-compliance.

3. Management Framework

Our proposed distributed requirements management framework provides traceability from the regulations that govern organizations to the decisions of actors who assign obligations to other actors and finally to the software requirements assigned to systems that refine those personnel obligations. Figure 1 shows the associations maintained in our framework among actors, systems and obligations. In assignment, also

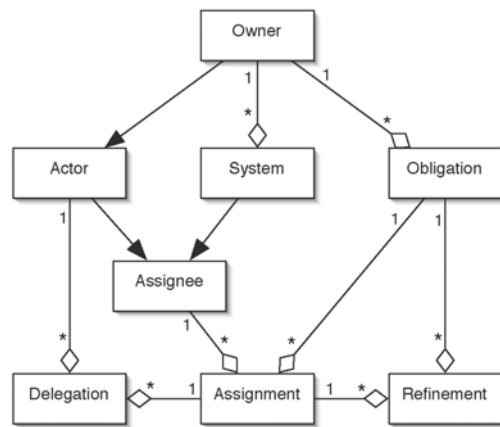


Figure 1: Framework Conceptual Model

called delegation, each assigned obligation has exactly one owner (an actor) who is ultimately responsible for satisfying the obligation. The owner may assign the obligation to other actors and the owner may further permit those actors to re-delegate the obligation to others. With regard to refinement, an actor who has

been assigned an obligation may choose to refine the obligation into other obligations, called *refinements*; satisfying these refinements contributes to satisfying the refined obligation.

Because the management framework seeks to mediate between personnel and information systems, we distinguish between two types of obligations: *responsibilities* that require a person to perform some action and *requirements* that require a system to have some property or perform some function. For all systems, there is exactly one *administrative obligation* that requires exactly one actor to be ultimately responsible for implementing the requirements assigned to those systems. In some situations this actor may be permitted to re-delegate her administrative obligation for a specific system to other actors. We discuss the important issue of authorizing delegation and refinement decisions to related work in Section 5.

We now define key terms to elucidate the primary elements of our management model.

Definition 1: Management Model

Let the set A consist of actors, the set S consist of systems and the set O consist of obligations.

The *assignment set* $AS \subseteq (A \cup S) \times O$ is a many-to-many relation mapping actors and systems to their assigned obligations; each actor or system may have multiple obligations and each obligation may be assigned to multiple actors or systems. Each act of delegation and refinement yields a new assignment.

The *delegation set* $DS \subseteq A \times AS$ is a one-to-many relation mapping actors (the delegator) to the obligations that they assign to other actors and systems (the delegatee). To be consistent, we require that all delegations yield a valid assignment, $\forall \langle a_i, a_j, o \rangle \in DS, \exists \langle a_j, o \rangle \in AS$. We also assume a permission framework is in place to ensure that each delegator is authorized to delegate obligations to the chosen delegatee.

The *ownership set* $N \subseteq A \times (O \cup S)$ is a many-to-one relation mapping actors to the obligations and systems they own; every obligation and system has exactly one owner. Each actor is solely responsible for monitoring the accountability of the obligations they own. For each system $s \in S$ with owner a in $\langle a, s \rangle \in N$, there is one administrative obligation o in $\langle a, o \rangle \in AS$ that requires the system owner a to satisfy all the assigned system requirements r in $\langle s, r \rangle \in AS$.

The *refinement set* $RS \subseteq AS \times O$ is a many-to-one relation mapping actors or systems and their assigned obligations to refinements (other obligations). These refinements are created by the actor who is assigned the obligation or by the system owner. A single obligation can be delegated to multiple actors, These actors may then refine this obligation into different specialized obligations in ways that satisfy the context of their daily operations. For an actor or system $a \in$

$(A \cup S)$ and an obligation $o \in O$, the set $\{r \mid \langle a, o, r \rangle \in RS\}$ is called a *refinement strategy*.

The *decision sequence* $\{d_1, d_2, \dots, d_n\} \subseteq (DS \cup RS)$ is derived by tracing an obligation through delegation and refinement decisions. A valid decision sequence is comprised of a series of delegation and refinement sequences: for some numbers m, n , a delegation sequence $\{\langle a_1, a_2, o \rangle, \langle a_2, a_3, o \rangle, \dots, \langle a_m, a_{m+1}, o \rangle\} \subseteq DS$ traces the obligation o from the delegator a_i to the delegatee a_{i+1} , and so on for $1 \leq i \leq m$; and a refinement sequence $\{\langle a, o_1, o_2 \rangle, \langle a, o_2, o_3 \rangle, \dots, \langle a, o_n, o_{n+1} \rangle\} \subseteq RS$ traces the obligation o_j assigned to the actor a to the refinement o_{j+1} , and so on for $1 \leq j \leq n$. A valid decision sequence begins with either a delegation or refinement decision $d_1 \in (DS \cup RS)$ and alternates between delegation and refinement sequences. To clarify how these sequences are connected, for some u, v, x, y : a refinement sequence follows a delegation sequence by $\{\dots, \langle a_u, a_v, o_x \rangle, \langle a_v, o_x, o_y \rangle, \dots\}$ where the delegatee a_v refines obligation o_x into obligation o_y ; and a delegation sequence follows a refinement sequence by $\{\dots, \langle a_u, o_x, o_y \rangle, \langle a_u, a_v, o_y \rangle, \dots\}$ where a delegator a_u refines an obligation o_x into an obligation o_y and delegates o_y to the actor a_v . We presently assume that no cycles exist in all sequences of delegations and refinements derivable from $(DS \cup RS)$; in practice, cycles can be identified and avoided. Refinement is complete if it is accountable, which we now discuss.

An obligation is *accountable* if a mechanism exists to verify that the obligation has been satisfied [12]. This mechanism may either be: (1) an oracle (e.g., executable program, hardware device) that returns true if and only if the obligation is achieved or maintained; or (2) the evaluation of a logical expression comprised of a conjunction of predicates. Each predicate denotes the satisfaction of obligations in a refinement strategy, each obligation of which must itself be accountable.

Consider, for example, the accountability of system requirements. Because *functional requirements* are testable by definition, some data or program exists called a *test case* that may be used to verify whether a system satisfies those requirements – hence, functional requirements are always accountable. *Non-functional requirements* are not testable in the same fashion, but they are accountable if they are refined into functional requirements. Therefore, testing non-functional requirements is tantamount to testing their refinements, assuming each refinement is either itself a functional requirement or another non-functional requirement that is refined into one or more other accountable requirements. Because delegation can transfer obligations from one actor to another, testing accountability in this distributed framework relies upon valid decision sequences.

Personnel responsibilities may also be accountable within this framework, if they are supported by software systems that retain sufficient information to evaluate those responsibilities. For example, the obligation “to only use a password that contains at least eight characters” is accountable by executing a program to check the length of a user’s personal password when it is set. However, the responsibility “to logout from a system when the system is no longer in use” is not accountable, as it is difficult to define the behavior of “in use” for all users, systems and applications.

Definition 2: Accountability

The *verification set* $VS \subseteq Boolean \times AS$ is a one-to-one mapping of Boolean predicates to assignments such that each predicate is true if and only if the assigned obligation is satisfied by the actor or system in the assignment. For a verification $\langle v, a, o \rangle \in VS$, the predicate v evaluates to either: (1) true if an *oracle* or *test case* decides the obligation is satisfied, or false otherwise; or (2) the logical conjunction of predicates assigned to some number of refinements: for the verification $\langle v, a, o \rangle \in VS$, let $v = v_0 \wedge v_1 \wedge \dots \wedge v_n$ such that $\langle v_i, a, o_i \rangle \in VS$ for all refinements $\langle a, o_i, o \rangle \in RS$. The expression must contain exactly those predicates for refinements o_i that are necessary and sufficient to satisfy the obligation o .

There will be situations in which only a human being can verify whether or not an obligation has been satisfied. In the HIPAA for example, there are situations in which medical information can only be shared with third-parties in the event of a medical emergency. Because medical emergencies can only be determined by appropriate individuals, an information system can at best receive this determination from a human but not verify the emergency itself. The framework can easily be extended to identify these oracles that receive these determinations. Moreover, auditing mechanisms can be put in-place to maintain logs of these oracles to identify misuse and for forensic analysis after an abuse of the system [14, 28].

4. Applying the Framework

We apply our framework from Section 3 to an example in which obligations (OB) from the HIPAA Security and Privacy Rules are delegated from upper management to their staff and later refined into software requirements (FR). At each delegation stage, an employee with specialized responsibility and technical expertise interprets his or her assigned obligations and refines and/or re-delegates these obligations, as needed. We illustrate this application by narrating the sequence of delegation and refinement decisions. At each state, we list the obligations

followed by the expressions in our model that record these decisions.

In addition to showing how our prototype framework would be applied to a real set of requirements taken from HIPAA, this example serves to show some of the complexity — and subtleties — present in legal requirements.

The Chief Security Officer (CSO) for a healthcare provider (a covered entity) is assigned the following obligation from the HIPAA Security Rule (SR) §164.308(a)(2):

OB₁: Identify the security official who is responsible for the development and implementation of the policies and procedures required by the HIPAA section 164 subpart C for the covered entity.

$$\langle SR, CSO, OB_1 \rangle \in DS$$

$$\langle CSO, OB_1 \rangle \in N$$

The CSO identifies the security official (SO) and delegates obligations from the Security Rule §164.302–§164.318 to the SO, including:

OB₂: From §164.312(a)(1): Allow access only to those persons or software programs that have been granted access rights as specified in §164.308(a)(4);

OB₃: From §164.308(a)(4): Authorize access to electronic PHI (Protected Health Information) that are consistent with the applicable requirements of subpart E of this part (e.g., the Privacy Rule); and

OB₄: §164.308(a)(4)(ii)(B): Grant access to electronic PHI, for example, through access to a workstation, transaction, program, process, or other mechanism.

$$\{\langle SR, SO, OB_2 \rangle, \langle SR, SO, OB_3 \rangle\} \subseteq DS$$

$$\{\langle SR, SO, OB_4 \rangle\} \subseteq DS$$

$$\{\langle SO, OB_2 \rangle, \langle SO, OB_3 \rangle, \langle SO, OB_4 \rangle\} \subseteq N$$

Furthermore, for those authorizations in the Privacy Rule that participate in a transaction and utilize an “electronic communications network,” the SO must also implement technical measures to:

OB₅: In §164.312(e)(1): Guard against unauthorized access to electronic PHI that is transmitted over an electronic communications network; and

OB₆: From §164.312(e)(1)(ii): Encrypt electronic protected health information whenever deemed appropriate.”

$$\{\langle SR, SO, OB_5 \rangle, \langle SR, SO, OB_6 \rangle\} \subseteq DS$$

$$\{\langle SO, OB_5 \rangle, \langle SO, OB_6 \rangle\} \subseteq N$$

Among the several authorizations in the HIPAA Privacy Rule (PR), we contrast the follow two authorizations to illustrate exceptions between rules;

for example, authorization OB_8 explicitly excludes reports of child abuse, whereas, OB_7 permits disclosing such reports:

OB₇: From §164.512(b)(1)(ii): Disclose PHI to a government authority authorized by law to receive reports of child abuse or neglect.

OB₈: From §164.512(c)(1)(ii): Except for reports of child abuse or neglect permitted by paragraph (b)(1)(ii), disclose PHI about an individual whom the covered entity reasonably believes to be a victim of abuse, neglect, or domestic violence to a government authority authorized by law to receive reports of such abuse, neglect, or domestic violence... to the extent the disclosure is expressly authorized by statute or regulation and either: (A) The covered entity, in the exercise of professional judgment, believes the disclosure is necessary to prevent serious harm to the individual or other potential victims; or (B) If the individual is unable to agree because of incapacity, a public official authorized to receive the report represents that the protected health information contained in the disclosure is not intended to be used against the individual and that an immediate law enforcement activity that depends upon the disclosure would be materially and adversely affected by waiting until the individual is able to agree to the disclosure.

$$\begin{aligned} \{\langle PR, SO, OB_7 \rangle, \langle PR, SO, OB_8 \rangle\} &\subseteq DS \\ \{\langle SO, OB_7 \rangle, \langle SO, OB_8 \rangle\} &\subseteq N \\ \{\langle SO, OB_3, OB_7 \rangle, \langle SO, OB_3, OB_8 \rangle\} &\subseteq RS \end{aligned}$$

The SO delegates these obligations to an Information System Architect (ISA) who is responsible for external disclosures of PHI to business and government associates, third parties, etc.

$$\begin{aligned} \{\langle SO, ISA, OB_2 \rangle, \langle SO, ISA, OB_4 \rangle\} &\subseteq DS \\ \{\langle SO, ISA, OB_5 \rangle, \langle SO, ISA, OB_6 \rangle\} &\subseteq DS \\ \{\langle SO, ISA, OB_7 \rangle, \langle SO, ISA, OB_8 \rangle\} &\subseteq DS \end{aligned}$$

The ISA refines these obligations into the following functional requirements:

FR₁: The system shall identify users by role: one role per law that (1) authorizes a user to receive reports of child abuse or neglect; (2) authorizes a government authority to receive reports of other abuse, neglect or domestic violence.

FR₂: The system shall identify data by subsets: one subset per law designating which PHI may be disclosed to users authorized to receive reports of abuse, neglect or domestic violence.

FR₃: The system shall record individually identifiable testimony from : (1) the user, an employee of the

covered entity, stating that they believe disclosing the PHI is necessary to prevent serious harm to the individual or other potential victims; or (2) the user receiving the PHI stating that the protected health information is not intended to be used against the individual and that an immediate law enforcement activity that depends on the PHI would be materially and adversely affected by waiting until the individual agrees to the disclosure.

FR₄: The system shall provide encrypted access to PHI identified in subsets (via FR_2) only to users identified by roles (via FR_1) only after receiving proper testimony (via FR_4)

$$\begin{aligned} \{\langle ISA, FR_1 \rangle, \langle ISA, FR_2 \rangle, \langle ISA, FR_3 \rangle\} &\subseteq N \\ \{\langle ISA, FR_4 \rangle\} &\subseteq N \\ \{\langle ISA, OB_2, FR_1 \rangle, \langle ISA, OB_4, FR_1 \rangle\} &\subseteq RS \\ \{\langle ISA, OB_6, FR_1 \rangle, \langle ISA, OB_5, FR_1 \rangle\} &\subseteq RS \\ \{\langle ISA, OB_7, FR_1 \rangle, \langle ISA, OB_7, FR_2 \rangle\} &\subseteq RS \\ \{\langle ISA, OB_8, FR_1 \rangle, \langle ISA, OB_8, FR_2 \rangle\} &\subseteq RS \\ \{\langle ISA, FR_4, FR_1 \rangle, \langle ISA, FR_4, FR_2 \rangle\} &\subseteq RS \\ \{\langle ISA, FR_4, FR_3 \rangle\} &\subseteq RS \end{aligned}$$

The ISA surveys existing systems within the covered entity and assigns the four requirements FR_1 - FR_4 to relevant systems. If existing systems are unable to satisfy any requirements, those requirements are assigned to a Software Engineer who will design, develop, test and deliver a new system or configuration to meet these requirements.

4.1 Compliance and Accountability

Heterogeneity between business practices in different organizations makes it difficult to develop a single, de-facto implementation of standards and regulations to achieve compliance. Consequently, auditors and external reviewers must certify that a set of business practices comply with a set of standards or regulations at a specific point in time: a process called *certification*. In addition, auditors must acquire real world evidence demonstrating that business practices continue to comply, either through random or continuous sampling of appropriate data: this acquisition is called an *audit*, and is performed as a form of *compliance monitoring*.

Using our framework, organizations can exhibit decision sequences that trace regulations to functional requirements. Each requirement is verified using one or more test cases to verify software systems. Auditors and external reviewers may certify that these sequences comply with law using a digital or cryptographic signature. The function $sign : K \times M \rightarrow S$ maps secret keys K and messages M to a unique cryptographic signature in S . The signature can be used to verify that a message, in this case the decision sequence, has not

changed and, furthermore, to digitally identify the auditor who certified the sequence. In follow-up reviews, decision sequences are re-certified using the history of verification predicates obtained during the policy and runtime requirements lifecycle.

Returning to the application of our framework in Section 4, the auditor identifies the following decision sequence $D = \{d_1, d_2, d_3, d_4\}$:

$$\begin{aligned} \{d_1 = \langle SR, SO, OB_3 \rangle, d_3 = \langle SO, ISA, OB_7 \rangle\} &\subseteq DS \\ \{d_2 = \langle SO, OB_2, OB_7 \rangle, d_4 = \langle ISA, OB_7, FR_1 \rangle\} &\subseteq RS \end{aligned}$$

The ISA exhibits the verification subset $V \subseteq VS$ consisting of verifications for all systems s in $\langle s, FR_1 \rangle \in AS$ and $\langle v, s, FR_1 \rangle \in V$. The auditors, deciding that the sequence D and verification set V are all necessary to comply with the regulations $\{OB_2, OB_3, OB_7\}$ at time t , uses their secret key k to provide the certificate $C = \langle \text{sign}(k, D + V + t), D, V, t \rangle$. In subsequent reviews at time $t' > t$, an auditor re-certifies these systems by verifying D and V against the previous certificate and, if the auditor accepts the result, the auditor will issue a new certificate updated for time t' . A series of these certificates for a single regulation at different time intervals is the compliance history for that regulation.

In the event of a violation of law or regulation, our framework provides a starting point for investigation. The violated obligation or requirement is identified, and then traced through to implementation and operation. Involved parties can also be identified from among all those associated with system development and operation. The result may be a faster, more agile response to addressing violations.

5. Related Work

In requirements engineering, traceability is a measure of quality that reduces the risk of not propagating relevant changes to artifacts throughout the development lifecycle. Maintaining traceability information can be time consuming [20, 36]; thus, some may view traceability as too costly. However, traceability reduces the risk of inconsistencies and ensures compliance of a resulting system with the Software Requirements Specification [1]. Within the context of regulatory compliance, traceability increases in significance and criticality because requirements are driven by multiple policies distributed throughout an organization. Introducing additional traceability within the context of compliance and accountability can affect the cost of these systems. However, the gains and benefits well outweigh the costs [36].

Research in system security has shifted focus to include broader issues in organizational management. We see evidence for this shift in both access control frameworks limited to permissions [8, 7, 32, 33] and more progressive approaches that also include

obligations and delegation [19, 29, 31, 34]. Permissions describe what people and systems are permitted to do, whereas obligations describe what people and systems are required to do given certain specific restrictions or constraints. In delegation, a person delegates her authority or responsibility to another person; the latter person acts on behalf of the former.

The proper authority to delegate obligations is important to ensure the standards and regulations can be properly refined into functional requirements. Barka and Sandhu consider modes of delegation including permanent and temporary delegation [8]. Revoking delegation authority would first require evaluating the impact on obligations that were refined and/or re-delegated. Bandmann et al. propose constrained delegation as means to moderate delegated authority in a distributed system [7]. Constrained delegation provides a means to control to whom a delegator can assign new obligations; this is important to prevent a delegator from obligating actors or systems that are beyond the delegator's scope of authority.

With regard to permissions, Oh and Sandhu use business units and organizational hierarchy to administer roles and permissions for users [32]. Park et al. attempt to align organizational structure with system structure to improve role-based access control implementations [33]. These approaches highlight the organizational need to conceptually align existing authorization frameworks with organizational structure.

Moffett and Sloman introduced the concept of policies and system objects [31], which they later realized in the Ponder language [15] to express authorizations and obligations for managing large networks. A deployment model for distributed network management was proposed using Ponder [19]. Minsky and Ungureanu introduce Law-governed Interactions (LGI) in which actors suffer penalties if they violate their obligations [29]. In LGI, actors subscribe to a shared controller that audits their behavior to detect non-compliance. Park and Sandhu propose $UCON_{ABC}$ to manage authorizations and obligations using conditions for digital rights management [34]. Each of these approaches shares common elements relevant to regulatory compliance, including the ability to express permissions, obligations and delegations and the means to audit compliance through obligations. Moffett identified the need for requirements in policy models [30]. For regulatory compliance, the refinement decisions of actors are also needed to completely trace from regulations to the requirements of systems that satisfy those regulations.

In requirements engineering, related work has focused on goal refinement and delegation [4, 5, 6, 17, 18, 23]. Goals describe desired states or actions performed by actors without specific consideration for normative positions (e.g., permissions,

recommendations and obligations.) Similar to obligations, goals are decomposed into sub-goals intended to achieve the original goal [5]. Darimont et al. describe the GRAIL tool [17] that implements the KAOS framework [18] for modeling goal refinement hierarchies using logical and/or relations and temporal logic. Regulatory compliance complicates the collaborative environment in which obligations are refined by personnel across an organization. Whereas the GRAIL tool goes far to address the rich semantics of goal refinement, it does not address the broader traceability issue where individual personnel decide how and when to refine obligations. Antón et al. show how non-functional requirements can be implemented through functional requirements [4]; a notion captured in our definition of refinement. Mylopoulos et al. describe the Secure Tropos framework that models ownership and delegation and defines obligation as “trust in execution” [23]. Similar to GRAIL, Secure Tropos provides a single-user perspective on goal refinement, whereas our compliance framework incorporates multiple viewpoints through distributed refinement. Bandara et al. propose applying goal refinement to policies using Event Calculus for temporal reasoning [6]. While goal refinement is not new, tracing refinement and delegation of obligations, together, in a distributed environment that supports audits and external reviews provides new opportunities to explore compliance-related issues.

Rees and Spafford describe an information security framework, named PFIREs, which combines policy assessment and review to mitigate risk in organizational security [35]. Although the PFIREs framework does not address policies as system objects, per se, the authors propose improving security by passing messages between personnel to communicate obligations and monitor compliance; these messages are necessary to implement a compliance framework such as ours.

The compliance auditing procedure we describe, using cryptographic signing, is similar to a current use case of the Tripwire tool [27] in government and financial settings: if the signature matches a saved signature, the code has not deviated from that previously certified.

Finally, several publications on policy, including the NIST Security Handbook [38], define policies as comprised of standards, guidelines and procedures. *Standards* are implementable obligations assigned to personnel and systems whereas *guidelines* are recommendations that may coincide with best practices for specific contexts. *Procedures* implement standards, often with a step-by-step description that is sufficient for other actors to reproduce the desired results. In our framework, these notions are complementary and can be easily supported. Guidelines require distinguishing

the modality of recommendations (should) from obligations (must, shall) and “a procedure” requires conditionally sequencing a set of refinements. Despite such extensions, it is important to emphasize that standards, guidelines and procedures typically originate with upper management, whereas our framework specifically supports middle and lower managers in their effort to contextualize organization-wide goals as obligations and requirements in a business unit.

6. Requirements for Tool Support

Building on our previous work in analyzing policies and regulations [2, 3, 9, 10, 11, 12, 13, 37], we propose requirements for a tool to support personnel from CSOs and business managers to software developers and system administrators. The work of CSOs is to organize and delegate the high-level goals that are (usually) relatively few in number and change infrequently, whereas, system administrators must respond to numerous system requirements that adapt systems to emerging business needs and security vulnerabilities. Furthermore, compliance officers and auditors need to evaluate the delegation and refinement decisions made by those personnel.

The following is a minimal set of software requirements (SR) for a tool that implements our framework. Following the requirements, we briefly describe how these requirements would interact with users as well as future extensions to the framework that would extend the capabilities of the tool.

1. EXPLORATION:

- SR₁:** The tool shall allow each user to view the “context” of an obligation. The context includes:
- (1) The delegator who assigned the obligation to another actor;
 - (2) The obligations, if any, that the assigned obligation refines (e.g., the overall goal the obligation is intended to achieve.);
 - (3) The refinements, if any, to the assigned obligation proposed by the delegatee.; and
 - (4) For a sequence of delegations and refinements, any certifications verified by digital signatures and the identities of signatory actors.
- SR₂:** The tool shall allow each user to separately view any obligations that are assigned to him by other actors.
- SR₃:** The tool shall allow each user with an administrative obligation to separately view the requirements assigned to the concerned systems.
- SR₄:** The tool shall allow each user to organize other users into groups such as business units or projects. Views may be restricted to only those users and obligations in a particular group.

2. ASSIGNMENT/ REFINEMENT:

- SR₅:** The tool shall allow users to assign obligations by user or group.
- SR₆:** The tool shall allow users to create rules that define the pre-conditions under which users are assigned obligations; the rules are automatically applied to users who will receive special indication, when viewing the obligation, that the obligation was assigned and by which rule.
- SR₇:** The tool shall allow users to refine their assigned obligations by creating new obligations and assigning those obligations to themselves or other actors.

3. CERTIFICATION:

- SR₈:** The tool shall allow users to certify a decision sequence using their secret key. The certificates (with digital signature) are used to indicate to other users the decision sequences approved by the users and to indicate that they have not changed since the certification date. The certifications are verified using the appropriate, known public keys.

A tool that meets these requirements will allow users to view their assigned obligations, refine these obligations into new obligations and delegate these obligations to others. In addition, compliance officers and auditors can expand obligations into two hierarchies that show delegations and refinements. These hierarchies allow compliance officers and auditors to investigate the decision chains introduced in Section 3 and discussed in Section 4.1. Furthermore, auditors can certify these chains and users can verify these certifications.

In addition, we believe users should be able to identify conflicts between obligations. These conflicts will require the owner or delegator of two or more obligations to assess the broader situation and prioritize these obligations. Another form of conflict or under-specification is an assigned obligation without pre-requisite permissions. Our framework can be extended to support user requests for permissions from those with the authority to delegate them. Using the tool, the delegators verify that the necessary obligations that justify the need for these permissions are assigned to the user. If those obligations are ever revoked, these permissions are considered for simultaneous revocation, ensuring overall consistency between policies and systems.

Furthermore, we foresee organizing standards and regulations in a custom set of plug-ins that contain pre-defined actors, obligations and refinements sufficient to align with an organization's business practices. These plug-ins should be developed by standards bodies and regulators and provided to organizations as part of the

traditional compliance package. An organization would align the actors provided by the plug-in with their own personnel, who would then refine their new obligations and align them with existing software systems. Alternatively, the users may identify the need for new systems that meet these new obligations.

Finally, software products could be distributed with their own requirements plug-ins. System administrators that bring a new product online would receive a list of requirements certified by the product developer, based on actual developer test cases. Whereas organizations that develop their own software assume their own liability for resulting software failures, these third-party certifications would establish third-party liability against explicit product requirements within the distributed management framework. Given a non-compliance event associated with a specific decision chain involving a third-party product, both the organization and third-party can determine the role the product played in the fault: the product may have failed to satisfy its requirements, or the requirements may have been inappropriately aligned with personnel obligations they could not reasonably satisfy.

7. Conclusion

In summary, we present a framework that combines delegation and refinement in a distributed system to capture the decisions that executives, managers, system administrators and developers make to achieve compliance with standards and regulations. The contributions in this work include a formal definition of decision chains that auditors can certify and review when determining compliance for an organization. Furthermore, we instantiate our framework using an example from the HIPAA Security and Privacy Rules and propose requirements for a tool to implement the framework with discussion of its use.

Our framework provides new structure that combines traceability and digital identity (via cryptographic signatures) with the domains of policy, law, software development and administration. Combined with our observations and discussions with practitioners, these preliminary results suggest the feasibility and effectiveness of the framework. However, relying solely on a tool carries specific risks. These risks must be augmented by meetings and ongoing discussions within organizations that use the tool [16]. The framework provides important traceability data on policy and requirements decisions that can be used to initiate and drive these discussions. Moreover, these discussions will likely yield interesting insight into additional interactions between requirements, policies and law that the framework does not currently capture. How can these interactions be described formally and integrated into the framework? Can we distill best practices for requirements engineers

and auditors from practical applications of the framework? We envision answering these questions by employing a light-weight, cross-platform tool based on the framework in future case studies.

Acknowledgements

This work was funded, in part, by NSF Cyber Trust grant #0430166: *Collaborative Research: Comprehensive Policy-Driven Framework for Online Privacy Protection: Integrating IT, Human, Legal and Economic Perspectives* and the Purdue Center for Education and Research in Information Assurance and Security (CERIAS).

References

- [1] A.I. Antón, R.A. Carter, A. Dagnino, J.H. Dempster and D.F. Siege. Deriving Goals from a Use Case Based Requirements Specification, *Requirements Engineering Journal*, Springer-Verlag, v. 6, pp. 63-73, May 2001
- [2] A.I. Antón, J.B. Earp, D. Bolchini, Q. He, C. Jensen and W. Stufflebeam. "The Lack of Clarity in Financial Privacy Policies and the Need for Standardization," *IEEE Security & Privacy*, 2(2), pp. 36-45, 2004.
- [3] A.I. Antón, J.B. Earp, M.W. Vail, N. Jain, C. Gheen and J.M. Frink. "HIPAA's Effect on Web Site Privacy Policies," *IEEE Security & Privacy*, pp. 45-52, 2007.
- [4] A.I. Antón, Goal Identification and Refinement in the Specification of Software-Based Information Systems. Ph.D. Thesis, Georgia Inst. of Tech., Atlanta, GA, USA, 1997.
- [5] A.I. Antón, W.M. McCracken, C. Potts. "Goal Decomposition and Scenario Analysis in Business Process Engineering." *Advanced Info. Sys. Engr., 6th Int'l Conf.*, Utrecht, Netherlands, pp. 94-104, 6-10, 1994.
- [6] A.K. Bandara, E.C. Lupu, J. Moffett, A. Russo, "A Goal-based Approach to Policy Refinement." *Policies for Dis. Sys. and Nets*, Yorktown Heights, NY, USA, pp. 229-239, 2004.
- [7] O. Bandmann, M. Dam, B.S. Firozabadi, "Constrained Delegation." *IEEE Symp. on Security Privacy*, pp. 131-140, 2002.
- [8] E. Barka, R. Sandhu, "Framework for Role-based Delegation Models." *16th Annual Conf. on Comp. Sec. Apps.*, pp. 168-176, 2000.
- [9] T.D. Breaux, A.I. Antón, "Deriving Semantic Models from Privacy Policies." *IEEE 6th Workshop on Policies for Dis. Sys. and Nets.*, Stockholm, Sweden, pp. 67-76, 2005.
- [10] T.D. Breaux, A.I. Antón, "Analyzing Goal Semantics for Rights, Permissions, and Obligations." *IEEE 13th Req'ts Engr. Conf.*, Paris, France, pp. 177-186, 2005.
- [11] T.D. Breaux, A.I. Antón, "Mining Rule Semantics to Understand Legislative Compliance." *ACM Workshop on Privacy in Elec. Soc.*, Alexandria, Virginia, USA, pp. 51-54, 2005.
- [12] T.D. Breaux, A.I. Antón, C-M. Karat, J. Karat, "Enforceability vs. Accountability in Electronic Policies." *IEEE 7th Workshop on Policies for Dis. Sys. and Nets.*, London, Ontario, pp. 227-230, 2005.
- [13] T.D. Breaux, M.W. Vail, A.I. Anton, "Towards Compliance: Extracting Rights and Obligations to Align Requirements and Regulations," *IEEE 14th Int'l Conf. Req'ts Engr.*, 2006, pp. 49-58.
- [14] F. Buchholz, E.H. Spafford, On the Role of File System Metadata in Digital Forensics, *Digital Investigation*, 1(4), pp. 298-309, 2004.
- [15] N. Damianou, N. Dulay, E. Lupu, M. Sloman, "The Ponder Policy Language." *Work. Policies for Dist. Sys. and Nets.*, Bristol, UK, pp. 29-31, 2001.
- [16] D. Damian, F. Lanubile, T. Mallardo, "The Role of Asynchronous Discussions in Increasing the Effectiveness of Remote Synchronous Requirements Negotiations," *Int'l Conf. Soft. Engr.*, Shanghai, China, pp. 917-920, 2006.
- [17] R. Darimont, E. Delor, P. Massonet, A. van Lamsweerde, "GRAIL/KAOS: An Environment for Goal-driven Requirements Engineering." *IEEE 19th Int'l Conf. Soft. Engr.*, Boston, MA, pp. 612-613, 2005.
- [18] A. Dardenne, A. van Lamsweerde, S. Fickas, "Goal-directed Requirements Acquisition." *Sci. Comp. Prog.*, v. 20, pp. 3-50.
- [19] N. Dulay, E. Lupu, M. Sloman, N. Damianou, "A Policy Deployment Model for the Ponder Language" *IEEE/IFIP Int'l Sym. on Intenerated Net. Mgmt.* Seattle, WA, USA, pp. 529-543, 2001.
- [20] Dömges, R., Pohl, K., Adapting Traceability Environments to Project-Specific Needs, *Comm. of the ACM*, 41(12), pp. 54-62, December 1998.
- [21] Ernst & Young, *Global Information Security Survey 2005: Report on the Widening Gap*, 2005.
- [22] Ernst and Young, *Global Information Security Survey*, 2006.
- [23] P. Giorgini, F. Massacci, J. Mylopoulos, N. Zannone, "Modeling Security Requirements through Ownership, Permission and Delegation." *13th IEEE Int'l Reqs. Engr. Conf.*, Paris, France, pp. 167-176, 2005.
- [24] "Standards for Privacy of Individually Identifiable Health Information." 45 CFR Part 160, Part 164 Subpart E. *Federal Register*, 68(34), Feb. 20, 2003, pp. 8334 - 8381.
- [25] "Standards for the Protection of Electronic Protected Health Information" 45 CFR Part 164, Subpart C. *Federal Register*, 68(34), Feb. 20, 2003, pp. 8334 - 8381.
- [26] D. Jackson, "Alloy: a lightweight object modeling notation," *ACM Trans. Soft. Engr. Meth.* 11(2): 256-290, 2002.
- [27] G. Kim and E.H. Spafford, "The Design and Implementation of Tripwire: A File System Integrity Checker" *2nd ACM Conf. on Computer and Comm. Sec.*; ACM Press; 1994.
- [28] B. Kuperman, *A Categorization of Computer Security Monitoring Systems and the Impact on the Design of Audit Sources*, PhD Thesis, Dept of Computer Sciences, Purdue University, 2004.
- [29] N.H. Minsky, V. Ungureanu, "Law-Governed Interaction: a Coordination and Control Mechanism for Heterogeneous Distributed Systems." *ACM Trans. on Soft. Engr. and Meth.*, 9(3), pp. 273-305, 2000.
- [30] J.D. Moffett, "Requirements and Policies." *Work.*

- Policies for Dist. Sys. and Nets.*, Bristol, U.K., 1999.
- [31] J.D. Moffett, M.S. Sloman, "The Representation of Policies as System Objects." *Conf. on Org'l Comp. Sys.*, Atlanta, Georgia, USA, pp.171-184, 1991.
 - [32] S. Oh, R. Sandhu. "Role Administration: A Model for Role Administration Using Organization Structure." *7th ACM Sym. on Access Control Models and Tech.*, Monterey, CA, USA, pp. 155-162, 2002.
 - [33] J.S. Park, K.P. Costello, T.M. Neven, J.A. Diosomito, "A Composite RBAC Approach for Large, Complex Organizations." *9th ACM Sym. On Access Control Models and Technologies*, Yorktown Heights, NY, USA, pp. 163-172, 2004.
 - [34] J.S. Park, R. Sandhu, "The UCON_{ABC} Usage Control Model" *ACM Trans. on Information and System Security*, 7(1), pp. 128-174, 2004.
 - [35] J. Rees, S. Bandyopadhyay, E.H. Spafford, "PFIREs: A Policy Framework for Information Security." *Comm. of the ACM*, 46(7), pp. 101-106, 2003.
 - [36] Ramesh, B. Factors Influencing Requirements Traceability Practice, *Comm. of the ACM*, 41(12), pp. 37-44, December 1998.
 - [37] E.H. Spafford and A.I. Antón. The Balance of Privacy and Security, To Appear: *Science and Technology in Society: From Biotechnology to the Internet*, ed. by Daniel Lee Kleinman, Blackwell Publishing, 2007.
 - [38] An Introduction to Computer Security: the NIST Security Handbook, NIST SP-800-12, Gaithersburg, MD, USA, 1995.