The Impact of Industry Constraints on Model-Driven Data Disclosure Controls

Tyrone Grandison[1], John Davis[2],

[1] Intelligent Information Systems, Healthcare Informatics, IBM Almaden Research Center,
650 Harry Road, San Jose, California 95120, USA.
tyroneg@us.ibm.com
[2] Chief Health Informatics Office, Office of Emerging Health Technologies
Veterans Health Administration, Department of Veteran Affairs, USA.
Mike.Davis@va.gov

Abstract. Healthcare data disclosure models, i.e. security and privacy models, have been created with the goal of meeting specific standard properties or principles, e.g. confidentiality, integrity, availability, limited disclosure, limited retention, limited use, etc. This approach has been widely accepted and used in many industries. However, examination of specific domain requirements leads to a re-evaluation of the operation of these controls and implies an uncomfortable realization, which is that the models may need to be augmented to take industry-specific factors into account at design time. In this paper, we propose a set of constraints that should be considered when designing security and privacy models.

Keywords: Data Disclosure, Security, Privacy, Constraints, Models, Industry, Healthcare.

1 Introduction

Data is the most important asset for a business or individual. Businesses need to protect data to ensure their competitive advantage, improve their bottom line and drive service delivery. Individuals, who provide this data, are primarily concerned with their data not being misused and with the level of control they have over the use and disclosure of their information. Unfortunately, the increasing number of data breaches reported by the media [1, 2] attests to the fact that these requirements are not being met.

Satisfying the needs of business and those of the client depend to a large extent, but not solely, on the disclosure technology in place to protect data and enable the creation of value. Data Disclosure Models (DDMs) are the designs for the specification, enforcement, validation and management of security and privacy policies and technology. These models are normally built upon a formal model of access privileges, computational frameworks, distributed computing notions and sometimes intuitive understanding of a process. A diverse range of models, which will be highlighted later in this paper, have been proposed and used extensively in

industry to understand and build corporate protection infrastructure. However, new evidence has emerged that the product of these models are being rendered ineffective in daily usage [3].

In [3], it was shown that the policies and mechanisms deployed to address the access concerns in a Norwegian healthcare organization were routinely circumvented because they were viewed as impediments to efficient job function. Similar observations [4-6] solidify the fact that this phenomenon is not the exception to the rule, but rather a general reality for existing systems. This raises some chilling concerns. The first is that the models used and policies in place are being rendered effectively useless, because they are perceived as road blocks and are thus often ignored and or bypassed. The second concern is that there seems to be an over-reliance on *retroactive*[1], exception-based mechanisms that create a secondary infrastructure, which over time may be a more realistic representation of system behavior and use. This over-reliance defines a state in which security and privacy controls are unable to detect breaches in real time. Both concerns imply a need to re-examine the creation of disclosure models to increase the use and effectiveness of the systems built from them.

The contribution of this paper is to 1) describe the considerations of existing enterprises that need to be included in model design, 2) introduce the current canonical set of data disclosure models, and 3) highlight how a sample model may be impacted by industry constraints.

2 The Basic Observation

The real world has many complexities and intricacies that make it a rich source of intriguing problems for researchers. Each enterprise has its own characteristic set of operating procedures, which is normally based on the domain that it belongs to. Data disclosure models, and their derived policies and technology components, are a way for computer scientists to work with abstractions that allow them to instantiate controls for a multitude of environments. For security controls, the conceptual underpinnings of the models currently in use involve using the company's business rules, its security policy, the range of purposes for which access is needed and the access decision evaluation algorithm or strategy to determine if information should be disclosed. However, these models often ignore a basic assumption.

Each industry or sector has at least one axiom that must be adhered to by any system or subsystem, computerized or not, that is involved in the production of its main deliverable. We will refer to this axiom as the *prime directive* for that industry.

A nurse who is tasked with delivering care to a ward of twenty patients, some of whom may not be in the care of a doctor associated with her, will rightly choose to override any disclosure controls governing patient data when the action is deemed warranted and or the situation is critical enough. The rationale is simple: *Nothing interferes with the delivery of care*. This is the *prime directive* for the healthcare

[1] In this context, *retroactive* mechanisms refer to methods and procedures that operate on information about incidents in the past.

sector. Thus, anything that contradicts this tenet is perceived as counter-productive and irrelevant.

Station employees at an entertainment company that are faced with a critical security or privacy violation in the midst of a money-generating task, such as a telethon or national commercial, must ensure that the revenue creation proceeds until the task completes. Thus, mitigation of a breach may be delayed in these circumstances. This highlights the *prime directive* for the entertainment industry: *Do not interrupt the cash flow*.

Stock market traders, analysts and other workers in the financial sector also have a similar principle that cannot be violated in practice. Breaches of the data disclosure technology that occur in the midst of trading and affect the stock ticker in any way have to be resolved after its effect will not be felt in the market. This may introduce a significant delay in violation resolution. However, it accommodates the *prime directive* for the finance industry: *Do not disturb the ticker*.

Based on the above examples, we purport that there exists at least one *prime directive* for each and every industry. The set of *prime directives* is not an infinite one. We also assert that it is possible to identify the collection of rules that embody these industry requirements, to characterize them and to develop profiles for security and privacy systems that can be easily deployed. This process would also be useful for data disclosure professionals who are changing industries and for corporations who wish to leverage these profiles to create targeted industry solutions.

Table 1 shows a general categorization of industry *prime directives* based on five dimensions: *Complete Knowledge*, *Emergency Access*, *Cash Flow*, *Conformity* and *Data Provider Empowerment*. *Complete Knowledge* refers to those industries that are grounded in the belief that either you know a fact or you don't. There is no gray area and ignorance is either non-existent or parameterized and quantified in their world view. An example of such an industry is Finance. *Emergency access* describes those industries where the normal modus operandi is to have a significant number of out-of-band scenarios occurring regularly, e.g. Healthcare. *Cash Flow* characterizes industries where each second of effort is expected to generate revenue, e.g. Web Commerce. *Conformity* represents the industries that are primarily focused on adherence of rules; where the rules could be regulation, system mandates or governmental imperatives, e.g. the Defense industry. *Data Provider Empowerment* describes those industries where data provider input is required before their data is processed, e.g. Medical Research. From the previous description, it becomes clear that an industry may fall into multiple dimensions and will thus be governed by several *prime directives*. For simplicity, Table 1 identifies the dominant dimension for each industry. Thus, Table 1 should be seen as a rough guide and it should be recognized that some of the industries listed may have secondary, and even tertiary, *prime directives*.

The characteristics column of Table 1 describes the industries, in that particular dimension, based on the cardinality of system states, the behavior when no credentials are available and the general attitude towards exceptions. These factors were chosen because 1) computer systems are state-based machines that can be in varying states of security and privacy compliance, 2) authentication is the entry point into current data disclosure systems, 3) *prime directives* normally represent exception states in most current systems.

For the finance industry, current systems assume that financial artifacts are either known or not. If one enters a brokerage house and does not know the account number or social security number associated with your account, then you are not allowed access to their system. As a general policy, no exceptions are allowed. In this case, given that there are two security states and that we know that security is not absolute, there is a crossover or inflexion point where the industry entity chooses a different paradigm: one that facilitates the dominant prime directive.

In the case of Web Commerce, security and privacy concerns are treated in the context of the impact magnitude of the highest priority risks. Thus, there are many security states. If credentials are not immediately known, then a credential recovery process can be executed and they can be retrieved and system access granted. Exceptions are allowed if they are deemed low-risk operations or transactions that do not affect business function.

Governing Concept	Industry	Characteristics	Directive
Complete Knowledge	▪ Finance ▪ Insurance	▪ Two discrete states, i.e. either you know or you don't know. ▪ No credentials → no access. ▪ No exceptions.	*Do not disturb the measure of revenue.*
Emergency Access	▪ Healthcare	▪ Many states ▪ No credentials → execute checks. ▪ Exceptions allowed.	*Do not interfere with service delivery.*
Cash Flow	▪ Media ▪ Web Commerce ▪ Communications ▪ Law ▪ Information Technology ▪ Advertising and Marketing	▪ Many states ▪ No credentials → execute recovery mechanisms. ▪ Exceptions that do not affect business operation allowed.	*Do not interrupt the flow of revenue.*
Conformity	▪ Defense	▪ Two states ▪ No credentials → no access. ▪ Exceptions allowed if a sanitization step is performed.	*Do not jeopardize the safety of people or systems.*
Data Provider Empowerment	▪ Research ▪ Education ▪ Services	▪ Many states. ▪ No credentials → execute checks. ▪ Exceptions allowed and require provider input.	*Do not engage in activity that conflicts with the rights or desires of the service consumer and or subject.*

Table 1. A Generalization of Industry *Prime Directives*

It should be noted that when an industry has multiple *directives*, the most general and least restrictive set of characteristics apply. For example, as Healthcare falls into

the *Emergency Access* and *Data Provider Empowerment* dimensions, the general characteristics that apply would be *Many States*, *No credentials → execute checks* and *Exceptions allowed*.

We leave it as future work to enumerate Table 1 with all possible industries. Evaluation of the architectural ramifications of the *prime directive* concept is also future work. It would be an interesting research project to ascertain if the prime directive is consistent with the natural laws of systems construction, i.e. the tendency for humans to build systems piecewise continuous with discontinuities at the endpoints, and normal system operation, which occurs in linear space. Answering the question: "Does the prime directive just simply mean that we at an extreme, where models are no longer valid?" would be a significant contribution to model-driven research.

The important implication of recognizing the existence of this phenomenon of a *prime directive* is that models need to be re-evaluated and re-designed. In addition, they may also need to be applied to, incorporated into and emphasize mechanisms that have received less attention in recent years than enforcement-oriented tools. This may translate into more focus on auditing, provenance, curation, watermarking and other *retroactive* technology. However, this basic observation is not the only one that affects model-driven systems.

3 Industry Constraints

Creating models and systems for industry requires a delicate mix of engineering knowledge, domain expertise and managed risk. Based on the authors' experiences working with systems in the healthcare, entertainment, finance, education and manufacturing sectors, we now enunciate the constraints that we believe should be an integral part of model-driven design for security and privacy technology.

1. **Thou shall obey the industry's prime directives.**
   The baseline rules for an industry that underpin the manufacture and delivery of its product or service and enable a positive bottom line must always be factored into model design.
2. **Thou shall synergize with the industry's surrounding ecosystem.**
   The integration of business considerations, social factors, legislative mandates and technology advances create a different effect for each industry in each country or sovereign entity. For example, the business climate, level of social awareness, state and federal laws and available technology in the United States create an adversarial framework, in which complementary security and privacy systems must be developed. This means that models need to account for this framework and include more checks and balances in order to protect all involved. These same factors in Europe lead to a collaborative framework that emphasizes openness and trust over accountability and redress.
3. **Thou shall be true to the system.**

Models should be created with the assumption that the deployed system expects the technology created to execute as if it is a proxy for the resource and or person. This implies that the model should reduce or eliminate the use of *trusted*[2] elements, because they may not always perform in the best interest of the system itself.

4. **Thou shall use resources sensibly.**
   Computing power, storage space, data flow re-directions and communication and IO bandwidth are issues that enterprises want minimized in order to reduce their costs. However, current models for data disclosure technology tend to introduce features that increase the use of these resources. The model designer needs to be cognizant of the fact that cost reduction or savings is an important factor that should be incorporated as much as possible. This is especially true in the case of minimizing data flow re-directions and reducing the points of exposure and failure.

5. **Thou shall be easy to manage**
   Security and privacy models should include the capability to maintain, update and remove the artifacts that they create with as little effort as possible. This capability could be encapsulated in separate models. However, they should also adhere to industry constraints.

6. **Thou shall integrate seamlessly into the workflow**
   Models should lead to systems that do not require the current workflow of an organization to change. Also, it should not be assumed that there is one typical workflow that describes all the operations of a particular field or domain. For example, in the healthcare sector, there isn't one complete workflow that captures the function of a family clinicians' office.

This enumeration of constraints should not be assumed to be exhaustive. It represents a starting point for model designers. We now focus on DDMs.

4 Data Disclosure Models

There is a large and diverse collection of data disclosure models [7-18], ranging from policy specification models to policy enforcement models to policy verification models, etc. Without loss of generality, we will restrict our discussion in this section to access control models. We now introduce five representative models.

4.1 Mandatory Access Control (MAC) Model

The MAC Model [7] enables the restriction of access to objects based on the sensitivity labels or classifications of the information contained in the objects and the

[2] In this context, *trusted* is used as articulated by the Trusted Computing Group (https://www.trustedcomputinggroup.org)

formal authorization or clearance of subjects (i.e. people, processes or devices) to access information of a particular sensitivity. The controls based on the MAC Model are normally enforced by the operating system or security kernel. For example, a typical control may be that the operating system will not convert a document with a secret label to one of a lower classification without a formal, well-documented declassification process.

4.2 Discretionary Access Control (DAC) Model

The DAC Model [7] enables the restriction of access to objects based on the identity of the subjects and or groups that they belong to. A subject has certain *discretionary* rights to pass on as permissions to other subjects. For example, Jackie may choose to give read access to his personal documents to members in the family group.

4.3 Policy-Based Access Control (PBAC) Model

The PBAC Model [8] enables the restriction of access privileges to authorized users based on the business classification of users and the organizational policies in place. For systems based on the PBAC model, the theoretical privileges, which are specified in policy, are often compared to actual privileges used and differences are automatically applied to inform better management of the system.

4.4 Role-Based Access Control (RBAC) Model

The RBAC Model [9] enables the restriction of system access to authorized users based on their roles. A role is essentially a collection of permissions, and all users receive permissions only through the roles to which they are assigned, or through roles they inherit through the role hierarchy. For example, the manager for the HR department is allowed to see his employees' benefit data.

4.5 Fine-Grained Access Control (FGAC) Model

The FGAC Model [10] enables the restriction of information disclosure based on policy and data-driven or subject-derived constraints. These constraints enable the access control system to exert control over the individual data elements based on data values, user consent or other arbitrary conditions. For example, Allan would like his demographic data disclosed to people at the bank doing account maintenance, but not to a bank-affiliated credit card offer company.

Each of the models presented above is indicative of a family of models, which adheres to the same core principles. It should also be noted that the PBAC, RBAC and FGAC models share some common, underlying concepts. A full exposition of data disclosure models is beyond the scope of this paper. The interested reader is

asked to peruse [11]. It should also be noted that each model is used to create the security policy and influence the evaluation strategy.

Traditionally, the MAC model has been associated with systems with a very high degree of robustness. It is assumed that the system's control mechanisms will be able to resist subversion and enforce an access control policy that is mandated by regulation. For this reason, we will discuss how the MAC Model is influenced by industry constraints in the next section.

5 Example of the Impact of Industry Constraints

Consider a simple American healthcare institution, AmeriHealth, with three subjects: John, Mary and Tyler, and three information types of interest: demographic data, medical data and DNA data. John and Tyler are doctors and Mary is a nurse. The system has three clearance levels (High, Medium, Low) and three sensitivity levels (Non-sensitive, Sensitive, Ultra-Sensitive). Fig.1 presents the formal authorizations for the subjects, the classifications for the information and the policy governing disclosure of this information.

Subject	Clearance
John	High
Mary	Medium
Tyler	High

Object	Sensitivity
Demographic	Non-sensitive
Medical	Sensitive
DNA	Ultra-sensitive

Policy

- Only subjects with a High Clearance can view Ultra-sensitive information
- Subjects with a Medium and Low Clearance can only view Non-sensitive information

Fig. 1. Application of the MAC Model to the AmeriHealth scenario

In emergency scenarios, Mary (nurse) will be required to see sensitive information for the duration of the emergency. However, given the model-driven disclosure system in place, Mary would be denied access. In this scenario, Mary would use the backdoor into the system in order to deliver care to her patient. If the model, and derived policy, had the capability to allow conditional access to different categories of data, then a backdoor would not be necessary for this type of situation. An example of such a system is presented in Fig.2.

In Fig.2, bypass mode provides a mechanism that allows the normal evaluation process in a data disclosure system to be circumvented when faced with a scenario that requires that a *prime directive* be invoked. However, this is merely one way of facilitating the *prime directive* in model-driven systems, such that it is a part of the data disclosure subsystem.

If a new patient, who happens to be an immigrant, enters the institution and is questioned about pre-existing conditions, would the system have any facility to ensure

that his preferences are complied with (industry constraints 2 and 3). Will his request that his information not be used for anything other than healthcare delivery be met? Will the system honor his requests and intentions? Currently, there is a high probability that the system will react in a manner that is inconsistent with the patients needs. This will be true for industries that are in an adversarial ecosystem, such as the USA. This implies that the notion of obligation, audit logs and redress should be more tightly integrated into either the primary MAC model or other associated models of the system.

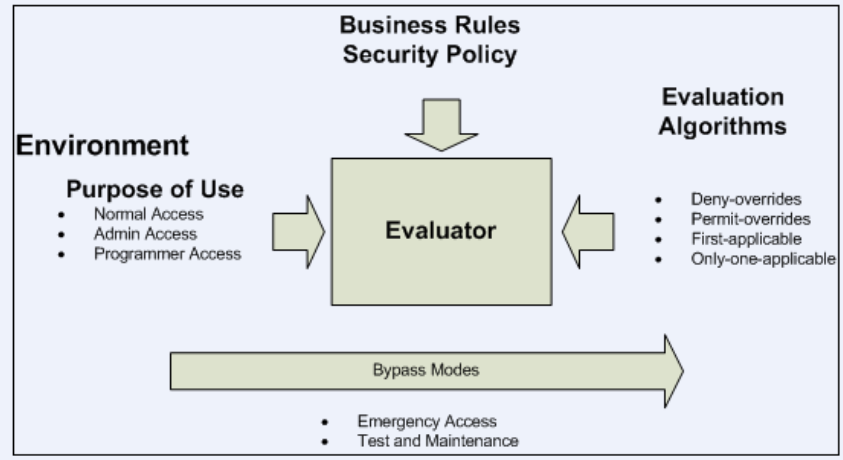

Fig. 2. Typical Data Disclosure Model with Bypass Mode

When the number of subjects increases or either the number of information categories gets larger or more granular, the space needed to store all this meta-information and the cycles taken to evaluate a particular security request become inhibitors to using the system (industry constraint 4). This implies that the MAC representation model may need to be optimized to handle systems with lots of authorizations, sensitivities and policy rules, e.g. by creating hierarchies of objects and subjects. This also highlights the fact that manageability may be a concern (industry constraint 5).

If a security system, based on the MAC model, where to be deployed at AmeriHealth, it would require that healthcare practitioners be cognizant of the their clearance level, the sensitivity of the information being accessed and the necessary steps required to make it available to others, inside and outside their team, at AmeriHealth. This imposes a cognitive demand on and shift in the performance of their duties. This extra burden, of becoming security-savvy, may represent a change to their workflow that they may not be comfortable with (industry constraint 6). Thus, such a system would have to be designed and engineered to reduce the security workload and workflow impact on the healthcare professionals.

In light of this very simple analysis, we can see that examining a data disclosure model in the face of the industry constraints specified in section 3, will most likely

uncover new features, thoughts and considerations that should be included in the model.

6 Conclusion

The central message of this paper is that data disclosure models need to be influenced more by the constraints of the environments that they will be used in. It is not feasible to create cross-industry models that can be applied in a multitude of geographies and for a wide range of clients. In this paper, we highlighted the fact that there is a trend towards circumventing model-driven data disclosure technology. This implies that the models that are being used to create these technologies need to be re-evaluated. Security and privacy models need to become more granular (i.e. targeted), if they are to lead to systems that are used more often than not.

We presented our initial observation in real world scenarios: For each industry, there is at least one prime directive that must be adhered to. This has been the primary cause for the circumvention of security and privacy systems. After making this basic observation, the other industry constraints that influence the acceptance and use of data disclosure technology were also articulated. Finally, we presented data disclosure models currently in use and described through a simple example how a model may change when viewed with industry constraints in mind.

References

1. Privacy Rights Clearinghouse: A Chronology of Data Breaches. http://www.privacyrights.org/ar/ChronDataBreaches.htm
2. Hanrahan, T., Fry, J.: A Million Good Reasons to Be Paranoid About Personal Data. Wall Street Journal, March,21, 2005. http://www.rmi.gsu.edu/rmi/faculty/klein/RMI_3500/Readings/Other/IDTheft_Internet.htm
3. Rostad, L., Edsburg, O.: A study of access control requirements for healthcare systems based on audit trails from access logs. In Proc. of the 2006 Annual Computer Security Applications Conference, Miami Beach, FL, USA, December 2006.
4. Pear, R.: Warnings over privacy of us health network. New York Times, February 18, 2007.
5. Break-glass an approach to granting emergency access to healthcare systems. http://www.nema.org/prod/med/security/upload/Break-Glass-Emergency Access to Healthcare Systems.pdf.
6. Tucci, L.: Electronic Medical Records At Risk of Being Hacked, Report Warns. CIO News. Sept 19, 2007. http://searchcio.techtarget.com/originalContent/0,289142,sid19_gci1273006,00.html?track=NL-48&ad=604676&asrc=EM_NLN_2220663&uid=4542080
7. United States Department of Defense: Trusted Computer System Evaluation Criteria. DoD Standard 5200.28-STD. December 1985.
8. Miller, D.V., Baldwin, R.W.: Access control by Boolean expression evaluation. Proceedings of the Fifth Annual Computer Security Applications Conference 1989. Pages 131-139.

9. Ferraiolo, D.F., Kuhn, D.R.: Role Based Access Control. Proceedings of the 15th National Computer Security Conference 1992: 554-563.
10. Agrawal, R., Bird, P., Grandison, T., Kieman, J., Logan, S., Rjaibi, W.: Extending Relational Database Systems to Automatically Enforce Privacy Policies. Proc. of the 21st Int'l Conf. on Data Engineering (ICDE 2005), Tokyo, Japan, April 2005.
11. McLean, J.: Security Models. in *Encyclopedia of Software Engineering* 2. New York: John Wiley & Sons, Inc. 1136–1145. 1994.
12. Levy, H.M.: Capability-Based Computer Systems, Digital Equipment Corporation. ISBN 0-932376-22-3. 1984
13. Lipton, R.J., Snyder, L.: A Linear Time Algorithm for Deciding Subject Security. Journal of the ACM 24 (3): 455-464. 1977
14. Biba, K. J.: Integrity Considerations for Secure Computer Systems. MTR-3153, The Mitre Corporation, April 1977.
15. Smith, R.: Multilevel security. in Hossein Bidgoli (ed.), Handbook of Information Security. Volume 3 - Threats, Vulnerabilities, Prevention, Detection and Management, New York: John Wiley. ISBN 0-471-64832-9. 2005.
16. Denning, D.E.: A lattice model of secure information flow. Communications of the ACM 19 (5): 236–243. 1976.
17. Sandhu, R.S.: Lattice-based access control models. IEEE Computer 26 (11): 9–19. 1993.
18. Bishop, M.: Computer security: art and science. Addison-Wesley. 2004