Information Risk in Financial Institutions: Field Study and Research Roadmap[1]

Sara Sinclair[1], Sean W. Smith[1], Stephanie Trudeau[1], M. Eric Johnson[2], Anthony Portera[2]

[1]Department of Computer Science, Dartmouth College
[2]Center for Digital Strategies, Tuck School of Business at Dartmouth
{sara.sinclair, sean.w.smith, stephanie.trudeau, m.eric.johnson}@dartmouth.edu

Abstract. Large financial firms with thousands of employees face many challenges ensuring workers have access to the right information, yet controlling access to unneeded data. We examine the problems of role lifecycle management and entitlement review processes in the context of large financial institutions. We describe observations from field study research in both retail and investment banks. We examine technologies to enable role and entitlement management and present a roadmap for future research.

Keywords: Access control, information risk, entitlement, provisioning, matrixed organizations, organizational complexity.

1. Introduction

Mergers in many industries have created very large global enterprises with thousands of employees, contractors, and partners scattered around the world. These massive collections of people lead to greater anonymity of the employees as they fade into the masses and can conceal actions that challenge modern security and controls [6]. Financial institutions provide a vivid example of this complexity. There are two major kinds of banks in the United States: retail banks and investment banks. Retail banks are the institutions that most of us are used to seeing on street corners, and advertise services like checking accounts and mortgages. Washington Mutual, Citizens Bank, and Sun Trust are examples of large retail banks. In contrast, most Americans will never directly interact with an investment bank. Investment banks serve companies and governments, raising capital by helping them issue stocks, bonds, and other securities. Investment banks often also engage in security trading services, hedge fund support, asset management for high net worth individuals, and analyst coverage of

securities traded in the public markets. Consolidation in the banking industry has created large investment bank firms including Goldman, Sachs & Co. with $37.7B in revenues and 26K employees and Lehman Bothers with $17.6B in revenues and 26K employees. There are also banks that have both retail and investment banking operations, such as Bank of America with revenues of $85B and 176K employees, and CitiGroup with $120B in revenue and 307K employees.

Such large financial institutions, critically dependant on information technology, must balance on a tightrope: information security compromise can lead to significant financial loss [5], but so can overly stringent security measures that prevent employees from getting their jobs done in a timely fashion. In this paper, we describe observations from field study research of both retail and investment banks. This research is the result of multiple visits to banks in each group. Moreover, members of our team spent three weeks embedded in the security group of one investment bank. During that time we also visited security teams in a set of New York City-based investment banks. This paper focuses primarily on internal information access controls and the risks of overaccess. We present field study observations along with a roadmap for future research.

2. Organizational Complexity Feeds Security Complexity

Banks are well-known for the intense expectations they place on their employees to meet the demands of clients. This intensity often results in significant organizational change and employee turnover. As professionals move within the organization, their information needs change, requiring rapid, precise computer systems that allow a great deal of flexibility, but without comprising security. It is no surprise that fraud prevention and confidentiality are major concerns of banks inside and outside of the firm. Confidentiality concerns focus on keeping client identities, strategies, intentions, and just about everything else held in confidence. Fraud spans issues from theft and misapplication to inappropriate use of confidential information.

Historically, opportunities for fraud may have been a greater concern in small firms where a handful of individuals administered the back office, each wearing many functional hats, and modern separation of duties and limited-access protocols were not feasible. As companies get larger, the intricacies of controls and needed permissions become more complex as both the number of systems and the granularity of access rights within enterprise-spanning infrastructures grow. In the large systems used by financial services, firms struggle to enable the right levels of access while restricting privileged information. Many users may not be aware of their access capabilities or the possible conflicts that their access could create.

A *toxic combination* is a conflict of system access permissions that allows a user to break the law, violate rules of ethics, damage customers' trust, or even create the appearance of impropriety. There are many ways for toxic combinations to occur. Sometimes it is a mistake of not terminating access following a promotion or transfer;

other times it is a fault of entitlement[2] design. An example of toxic combinations occurring from promotion could be as seemingly innocuous as an accounts payable clerk retaining the access to write checks once he has been promoted so he can assist with that task at busy times. If his new job allows him to go back to edit and even delete check writing records, he has the opportunity to steal money while circumventing traditional checks and balances. A design flaw example would be a trader in a commercial bank having access to see holdings of the accounts for clients she manages, as well as those of other trader's clients. The trader's access could be used to counter the aggressive positions of her non-direct client to the enrichment of herself and others, which is not only unethical, but also highly illegal.

Over-entitlement is also risky. It may not seem problematic for employees to have access to systems they never use or are unaware of. However, such access introduces risk. The root of the problem is that unnecessary or uncontrolled access can lead to unintended data editing, accidental disclosure, or internal misuse. That is why Sarbanes-Oxley auditors will flag unnecessary access as a weakness. The large investment banks have thousands of information systems and millions of different entailments resulting in extreme complexity for new or transfer employees to get the permissions they need. A common solution to this problem is that a new employee or a transferred employee will generate a request to simply copy the entitlements of the most well-endowed person in the department to the new individual, frequently leading over-entitlement for the new individual's needs.

From an employee's point of view, increased portability and accessibility of information facilitates productivity. Employee turnover has always been a concern for information leakage, but, as information becomes more fluid and more easily accessed over public networks, control of voicemail, remote email, PDA email, and home use of corporate files becomes more important. For example, a few years ago, one large bank's IT security group came to the realization that much of their important data lived in Excel spreadsheets, rather than the large, secure proprietary systems one might have expected (and hoped). This realization led them to prioritize the purchase and implementation of additional software and controls to further limit the movement and access of these files.

In addition to the large, multi-function corporate systems, companies are also shifting management strategies and human resources faster than ever before. In the investment banking space it is quite common for people to move between internal organizations and be transferred across information boundaries. The frequent shifting of staff may result in information users collecting system entitlements over time if the system access is not actively managed, resulting in a toxic combination of privileges.

Innovations in organizational structure also make security more difficult [7] as far as approving, monitoring and terminating information access are concerned. Most people are familiar with static hierarchical organizations in which everyone has a boss, who reports to a higher-level boss, and organizations have defined lines of responsibility as depicted in Figure 1. Team-based and matrix-structured organizations [1] are becoming more and more prevalent in professional society and especially in

---

[2] An *entitlement* is a resource that a person can be authorized to access in a certain way; for example, "opening case files" might be an entitlement for application X (we could also call it a "privilege" or "permission").

the financial services industry. In a matrix organization [3] employees may report to many "bosses," each for a set of different projects or activities, as also depicted in Figure 1. Some entitlement technologists have called these intertwining multiple hierarchies *polyarchies*. The polyarchy is simply a group of non-linear reporting relationships that come together to make an organizational matrix.

As more organizations take on a matrix structure, it becomes less evident who reports to whom and who is responsible for permitting and terminating data access. Employees may no longer have direct managers, as roles such as functional manager, group manager, engagement manager, review manager, and co-manager have become prevalent in financial institutions, professional service firms, and corporations. Additionally, the need for easy transfer from one role to another within a company is part of the organizational design that accentuates the complexity of securing information. Questions like, "Who owns this data or file?" and, "Who should approve access?" become common as data stores accumulate. The natural trap is for IT to give people access to whatever they ask for without an appropriate approval process; this often happens in cases in which there is no credible manager and IT is not able to track when the entitlement is no longer needed, or identify when it conflicts with new responsibilities. We have heard IT professionals jest that one can track a person's career path by examining their system entitlements, which seem to rarely be adjusted down in advancement or transfer. We have been cited examples in which 50-90% of the individuals with access to particular data store also have legacy access to information that they no longer need.

One critical goal of information security in the financial industry is to get the right information to the right people with maximum efficiency. Often, this involves the laborious task of individually granting access rights and entitlements of specific program applications to each user. A growing desire from security professionals within the industry is to create automatic rules and roles that can identify user attributes and access needs so as to automatically perform entitlement and potentially anticipate future permission needs. This goal is not dissimilar to that of sophisticated on-line retailers and content providers. The recent buzz among the management for websites like Yahoo.com and Facebook.com is that they aspire to create a "segment of one." Essentially, they would like to perfectly tailor the content a user receives to precisely meet the needs of that user based on historical and current activities. Given rapidly evolving content availability and the vast amount of traceable data a user can generate, perfect individual tailoring is nearly impossible without dedicating vast amounts of human resources to constantly code new rules or artificial intelligence that can learn and predict trends and personalities. While the "segment of one" may be far off, the institution of automated rules to guide information access may not be far from reality for investment banks. These corporations' desire for automatic permission tailoring is similar to customized content delivery, but the risks of getting it wrong are higher both in productivity downtime and unintended access. We will further discuss this issue after we examine some key security strategies employed by several investment banks.

Fig. 1. Privileging in traditional hierarchical corporate structures (left) vs. in dynamically, "matrixed" organizations (right). An arrow represents a supervising relationship (directed graph). We note that on the left, each person has exactly one direct supervisor, whereas on the right, each may have two or more.

3. Security Strategies

As part of our research, we interviewed IT professionals in a collection of large Wall Street financial institutions. We found one firm to be highly progressive, whose security professionals see the ambiguity of their matrix organizational design as an opportunity to "get into the business" and build collaborative bridges between IT and revenue-generating operations, both adding value to solutions to improve operations and generating a more secure environment through understanding the issues surrounding the business. This practice of getting "into the business" requires a highly competent IT staff, a self-starter culture that fosters collaboration and self-sacrifice for the good of the organization, and the tolerance for the additional investment in relationships. While expensive, it seems to result in bleeding edge innovation and preventative results that seem to pay big dividends in the complex, fast paced world of financial markets.

   In this particular financial organization, security strategy is tied to its culture. This culture is strongly oriented toward self-starters who work in teams with the goal of enriching the firm immediately through profitable activities, and over the long term by building the firm's reputation. This view that reputation is paramount in importance to the long-term fortune of the firm and individuals motivates the firm leaders to do whatever is conceivably possible to protect the IT resources of the firm within the bounds of productivity requirements. The urgency to protect the firm results in a few key strategic positions.

   Willingness to Buy or Build New Applications: It sounds simple to say, "If you cannot buy something, you can just build your own." However, this is not a common approach outside the upper echelon of the financial services industry, and even there, building custom solutions can be seen as too laborious or expensive to make it worth the firm's while. This particular progressive firm is truly agnostic on the buy-versus-build decision in its opinions on headache and firm resources. It feels that the security of its specialized systems and data is too important to risk not having the right system attributes. To this end, it has an army of employee programmers and contractors at the ready to build what it needs for both security and business applications. One staffer

notes that they are "totally fearless" when it comes to building a solution if they cannot find a commercial one that meets their needs. They are also not afraid to ask a vendor to sell them part of a program (and then build it into existing or new applications currently used by the firm), or to request that a vendor heavily modify an application. One employee stated that this determination to find the right solution, even if it doesn't currently exist on the marketplace, "teaches them a lot" about what they need and what to look for in future solutions.

At this firm, an initiative to build a new program for the firm is not taken lightly. Employees are highly intolerant of half-baked applications (whether created internally or by vendors) that unduly slow or block the flow of business for a moment. The decision to build an application is a commitment to see it through to its operational completeness.

To contrast this firm's perspective, another firm we spoke with had the capability to and does build its own systems for productivity and security functions, but we detected its bias to purchase software or work with an outside vendor who will develop a commercial application that includes their needs. Their perspective is based on the benefits of financial efficiency of using outside vendors and appreciates the support that comes with purchasing software.

Financial Freedom to Explore Options: The delegation culture that results from management's expectation that all employees of this particular firm are self-starters can allow for great amounts of freedom to explore new concepts in information security without burdensome oversight. This freedom is provided via financial funding and an absence of day-to-day accountability for personal actions within the project management level of the security group. The security group is *expected* to keep the firm safe. As long as security breaches do not occur or, more importantly, that nothing shows up in the *Wall Street Journal* indicating a confidentiality or fraud event, the group does not fall under scrutiny and feels free to work on whatever initiative seems important to them, even if the results are not immediately visible.

The Chief Information Security Officer expressed his constraints well: "Funding and resources have never been much of an issue. It is a question of control and flexibility. The hardest part is to achieve balance from the client side." The control of the information resource and flexibility for the business to get the job done are of paramount importance. Each one must be satisfied while not overpowering the other.

This freedom allowed this firm to develop innovations in information rights auditing for key applications years before Sarbanes-Oxley came into existence and regulators began penalizing other investment firms for their lack of foresight. This firm's institution of the SOX-type tools was not motivated out of the expectation that Congress would institute new rules, but that the firm and the industry needed this type of controls.

Open Vendor Relations: Despite its willingness to build whatever it needs, this firm is a big buyer of vendor-created applications. However, its size and clout within the industry leads to relationships that are generally more than the standard arms-length buyer-seller relationship. These relationship differences are derived from this firm's objective in getting precisely the right program attributes and willingness to collaborate. It is not uncommon for this firm to ask technology vendors to customize their products and make them go through months of, and sometimes years of, proof-of-concept exercises.

This firm also sees itself as an innovator within the industry. Management often shares the program attributes that it has created generally within its internal systems with vendors who approach them with "competing" products. They see this as both an opportunity to help the vendor create a product that may be useful to them in the future and an opportunity to give vendors insights that may help others in the industry.

Innovative Security Group Structure: Many investment banks separate the responsibility of monitoring security needs and generating security policies to a business risk organization staffed with business analysts who frequently work with a separate IT security organization that is responsible for the development and management of security applications. At this firm, however, there is no separation of the business risk organization and IT risk organization. The security innovators are the security creators. Many of the professionals who join the team come in with backgrounds in various technical and non-technical fields—but not security. They learn security on the job and are expected to be conversant in both the IT needs and the business activities of the firm. When they discover a need to mitigate risk they do not give the task to another organization they find their own solution or directly manage its development with the programming resources within the firm. They have created an environment where policy creation is not the end of one person's job that is handed off to another, but the beginning of a development cycle where the policy developer is also the solution creator who must prove hypotheses and get firm buy-in for the changes he spurs in the organization.

As a security application or tool is created, the team makes sure it can be "commoditized." They use the term commodity in this sense to indicate that they do not want the usage of the tool to be so complex that it cannot be handed to someone else. In fact, the security team does not want to build a security empire. They want to develop tools and spin them out to fall under the responsibility of other existing or new organizations that report outside the information security group; one employee stated that they see themselves as a "factory solution" that creates and distributes technology and processes to the rest of the firm. The security team retains some control or input into the strategy, architecture, and budget of the spin-out, but does not want to be part of the day-to-day administration of security, just the innovation.

This structure may not be fully feasible to other companies, but the practices of eliminating boundaries by combining functions and raising the bar on security professionals to see themselves as part of the business could be helpful to many organizations.

4. Managing Complexity in Security: A Field Study

Thus far, we have examined characteristics that can affect information security in financial institutions; we know that large firms are particularly challenged by having thousands of employees performing diverse tasks across the globe, and also because these overarching organizations are usually composed of many independent business units facing interdependent security risks [8,9]. Furthermore, we have examined elements of the culture at this specific firm and customs that directly impact the development and deployment of security technology, noting the "build or buy" policy, the

financial and operational freedom provided for the establishment of secure systems, the company's relationship with software vendors, and—perhaps most significantly—the unique structure and philosophy of the technology risk group.

To better understand the general processes that a financial institution goes through in solving information risk problems, as well as the ways in which the above-mentioned elements affect these processes at this firm in particular, we have considered a number of real-world issues in detail. Specifically, we have researched the way in which entitlement and provisioning is changing at this firm, and what (rule and role-based) technologies are being used to adapt existing systems and infrastructure to accommodate needed control and review capabilities.

4.1 Entitlement Review

During our research visit, various teams were in the process of finishing a comprehensive entitlement review of all SOX-relevant applications. The purpose of the review was to verify that each employee of this firm who has access to these applications has an appropriate set of entitlements, neither so few that they can't complete their jobs, nor so many that they pose an information risk. In many lights the review was a success: most employees' active entitlement set was reduced 30-50%. Yet, although the participation was high enough to satisfy audit requirements, the review also saw a low level of buy-in from certain business units (even when we account for the business units that were exempt because they already had an entitlement review process in place); given these two statistics, it is difficult for this firm to estimate the impact that the review had on the risk posed to the company. As one staffer told us, they cannot know if reducing Alice's entitlements twice as much as Bob's indicates that Alice's risk was reduced twice as much as her counterpart, nor can we accurately estimate how business unit A's 100% participation compares to business unit B's 50%. The technology risk team deals with this mix of hard statistics and nebulous interpretation all the time, which makes measuring the success of their efforts (and thus the value to the firm) very difficult.

Employees were asked during the entitlement review process to undergo a self-review, during which they could voluntarily give up entitlements they knew they did not need. This phase of the review was more productive than the staff expected, for employees collectively reduced their entitlement to applications (not just individual entitlements within an application) by 15%. Technologists said that employees "just didn't want to worry about" having access to applications they didn't need. This could be in part motivated by the fact that employees' managers were to perform a review of each employee directly following the self-review; an employee knew that they could only benefit in their manager's eyes by giving up entitlements. Given other observations, we feel that this voluntary de-privileging was perhaps more subtly motivated by the larger culture toward doing what's best for the firm. If employees genuinely have the best interests of the company at heart, and understand the risk over-entitlement poses, they would choose to give up as much as they could.

Unfortunately, this enthusiasm for de-provisioning during the review resulted in under-provisioning for many employees. In some cases, this was because employees could not correlate the entitlement description in the review to an entitlement they ac-

tually use. This was often because the human-readable entitlement description required of each application owner was not sufficiently clear, or was missing, but sometimes attributable to the users simply not knowing that clicking on the green button on their computer was really launching application ABC. In other cases, employees became under-provisioned because there were hidden entitlement dependencies, i.e., it was not clear that access to application ABC was essential to important work in application XYZ.

Why was the entitlement review process difficult? Shouldn't it be straightforward to list the number of applications, list the different privileges within them, and evaluate the mapping of those privileges to users? Shouldn't it be clear which entitlements users need by reading their job description and project assignments? This is, in fact, an area in which the size and complexity of the company inherently complicate the task at hand. The vast majority of applications are used and maintained by specific business units; there are thousands of applications spread throughout the firm, and new ones are constantly being deployed. Each application has its own notion of entitlements, its own entitlement descriptions, and required its own mental model of how to map human privileges to data access (and there are tens of thousands of humans who might be users for each system). This results in each user having hundreds of individual privileges, and supervisors being forced to review potentially thousands of individual entitlements (which, again, are not necessarily represented in a very human-understandable way).

The biggest challenge isn't the massive number of entitlements or users, however, but the highly dynamic nature of employees and organizational structure within the firm. The matrixed environment is hard to evaluate, and even harder when an individual's manager and entitlement needs are in constant flux (even though this dynamicism is a source of strength in the business world). The only way staffers were able to complete the entitlement review is because they took a "snapshot" of the entitlement systems, and only re-evaluated this snapshot a few times during the review process. During a few months of the review, one business group of 3,000 people witnessed 1,000 changes to organizational structure; in the space of a few weeks, 158 users in another group had changed job positions. If the process took into account all of these changes as they were happening, employees conducting the review would be so busy updating the picture of organizational structure that they would have no time to actually review entitlements. Conducting an entitlement review annually would be much easier if there were a persistent and up-to-date picture of the organizational structure and entitlements throughout the year.
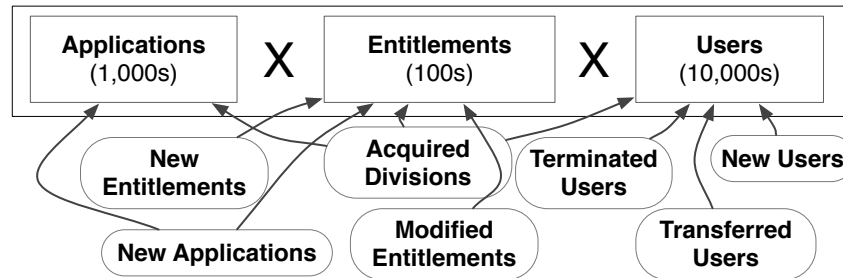
Fig. 2. Complexity and dynamics in entitlement systems. The number of applications, entitlements, and users make entitlement a large-scale problem, and the number of daily modifications to each of these sets makes it a fast-moving target. This algorithm does not consider any attribute of users other than the entitlements they currently have.

4.2 Using Roles for Structure and Entitlements

Many corporations are looking for the best way to manage entitlements in this kind of large, complex, and dynamic environment. Employees of some companies we have interviewed doubt the feasibility of any kind of centralized management of entitlements; they feel that the space and diversity of data is simply too massive. However, we found in the firm we observed not only a willingness to consider such technologies, but deployment of one already in progress. The team of technologists behind this implementation is driven by a vision of a *role-based* system that maps entitlements to classes of users via *rules*.

Organizations such as the NIST have been championing Role-Based Access Control (RBAC) schemes [11, 12] since the mid-nineties, and there exist a number of frameworks to implement access control using this model [10]. RBAC traditionally affiliates business-level functions, or roles, with sets of permissions on a given system; this grouping of entitlements facilitates provisioning (e.g., Alice is a desk manager, so we give her access to the same kinds of resources as other desk managers), as well as entitlement review (e.g., all desk managers should have the set "S" of entitlements) and modification to existing entitlement groups (e.g., we want to give all desk managers access to system XYZ when it comes online). In theory, RBAC schemes allow us to segregate the massive numbers of employees and entitlements into distinct groups that are easier to manage. However, the size and complexity of large banks make role-based systems challenging. At one very large retail bank that we interviewed, the CISO had recently completed an RBAC project creating 11,000 roles across the firm to control access to its nearly 22,000 applications. Developing the roles took a team two years and the ongoing review process was expected to be significant.

To study the feasibility of deploying a role-based system, one team at another firm we observed set out to estimate the number of roles that would naturally fall out of existing applications; they found that the number of simple roles within the company greatly outnumbered employees. This is a reasonable result given the number of applications and entitlements, but for roles to bring value to the company, the system

must group users into more manageable units, not just add a layer of complexity and abstraction. A role-based system can clearly mitigate some forms of complexity, but offers also its own set of management challenges in defining appropriate roles. Furthermore, we can only assume that the dynamicism inherent to organizational structure within an investment firm will lead to highly dynamic roles as well; as jobs and applications evolve, so will the logical groupings of employees doing those jobs and using those applications.

In response to these observations, the team who is integrating role-based systems into the corporate infrastructure is focusing on role management, and on the *life-cycle events* that define a role and determine which users are assigned to it. For example, a new employee joining the firm is an event that triggers assignment of appropriate role(s) to her: when that employee changes divisions, she may lose some roles and gain others; when she leaves the firm, she should be removed from membership of all roles. The series of events starting with an employee's hiring and ending with proper de-provisioning upon her leaving reflects the grand vision of role management within this firm; through constant event-based roles updating, it should be possible to get a clear picture of the entitlements currently assigned, as well as easily change entitlements for classes of users when necessary.

Thus far, staffers have identified some key characteristics of a system that would meet this grand vision. (They have also contracted to use a commercial product to help with role management, which we discuss later in this document.) Most importantly, they believe that role management will consist less of technological solutions, and more of business processes. This is in line with prior work that suggests many security problems are due to a lack of understanding of human systems [13], or to technology that does not appropriately model the needs of those human systems [14]. Parallel mechanisms already exist in entitlement and provisioning systems; there are well-defined ways in which business units communicate with human resources managers, or in which application users can request new access privileges from a particular application owner. Integrating roles and role management into their infrastructure will require discovering the correct business processes to tap into, and inventing new business processes where appropriate ones do not exist. Security technologists at the firm will be heavily involved in the initial deployment of roles, but once they are established, business-focused employees should run these systems; to ensure that they maintain up-to-date in the dynamic environment, they must be managed by those already effecting dynamicism within each business unit.

This process-based vision of role management is different than the traditional research-based idea of roles. Indeed, this firm's definition of a "role" as a persistent group of entitlements with a name and correspondence to business role is different from most instantiations of the RBAC model; in the latter, an individual user often has many roles, and only takes on one at a time. The traditional RBAC model strives to make sure that a user only has access to the set of entitlements with the least privilege to accomplish a given task; in this way, if a user's access is compromised, the number of resources affected is minimized. The firm's view, in which users always employ the roles assigned to them at all times, is much more similar to "group"-based access control schemes. However, as one staffer noted, the types of groups they need are much more sophisticated than current implementations (such as Active Directory).

This is an area of authentication theory we hope to explore more with this firm, and share with the research community at large.


Fig. 3. A set of users may have entitlements for resources X, Y, and/or Z (b); here small shapes represent users, and shade is determined by the set of entitlements the user has. A "blind" clustering algorithm (a) groups users into roles based on the sets of entitlements they have. Here, users with entitlement X (light grey) are grouped into Role A; those with Y and Z (dark grey) are in Role B, and those with all three entitlements (black) are assigned to Role C.

Although this firm recognizes the need for effective role management mechanisms, they also understand that they must establish roles before they can manage them. However, because they hope to use roles to get a better grasp on the different types of resources in the firm, we find a bit of a chicken-and-the-egg problem: how do we define roles as a function of resource entitlements, when we're planning on using roles to help understand what our resource entitlements are? Given the scale of the problem, the CISO scoffed at the idea of manually grouping people into roles.

The commercial product that the firm has licensed features a "role discovery" function, which allows system owners to mine entitlement information for pre-existing clusters of privileges that might be incorporated into roles. This approach is very attractive, especially in a large matrixed organization; it promises to automate what would be a very tedious manual process. However, we joined our collaborators at the firm in questioning the scope and the value of role discovery, and uncovered a line of inquiry on the nature of roles and their evolution. Because role discovery (as implemented in this case) operates by blindly clustering groups of people (Figure 3) with similar entitlements, how can we be sure that the resulting roles can map meaningfully to existing business roles (Figure 4)? If there is no mapping to business roles, what happens when we decide to add or remove entitlements within specific roles? A role-based system's ability to apply entitlement changes to large groups of people is an attractive feature; if our clusters of users are poorly defined, we might find the need to split, merge, or otherwise redefine role membership from the very beginning.

Fig. 4. Alternate clustering algorithms from that presented in Figure 3. Depending on the nature of additional information, such as job status (temporary or permanent, represented here by shape), better definitions of roles may be dependent on more than just entitlement data.

Of course, as the firm evolves, we expect there to be some redefinition of roles. For example, one business unit may divide its accounts into two sets, and determine that there must be a Chinese wall between them. Before the division, perhaps there was a Role A that encompassed the entire unit. After the division, should the group be split into multiple roles: X for those who have worked with one set of accounts, Y for those who have seen the other set, and Z for those who have not yet stepped off the Chinese wall? Or should the role acquire an attribute that allows the system to identify subgroups, i.e., attributes A_X and A_Y, which can overlap with other characteristics in the same role? In subtler cases, might we see gradual "drift" in roles, in which a group or subgroup diverges over time from the original role definition?

These possibilities indicate that role discovery might be a useful tool in first implementing a role-based system, but that there may be elements of the nature of roles that prevent it from being useful in already running systems. (Then again, if the role discovery algorithm were modified to not just blindly cluster, but also take in other information that can influence the grouping of people into roles, it could prove to be more useful in the long term.) These questions also reflect a certain level of ignorance on our part about the nature of roles; currently, we cannot predict how they will evolve once implemented. In a larger context, this problem represents a larger inability to predict the effects technical changes in a system that involves both humans and computers will have; we can simulate changes in networks or data centers, but human systems add new and puzzling factors. We know that the firm we observed is considering these kinds of questions in its plan for role deployment, and hope that through deployment they will be able find some answers questions; we hope to extend our collaboration to include research in this area.

4.3 Role-Based System Technology and Deployment

Different parties in research and industry have developed different notions of "roles," depending on their needs. We have discussed the importance of role life-cycle management in deployment of a role-based system, but we must also consider the technical elements that will make this deployment possible. Perhaps even more relevant to our goals of understanding the security technology deployment in the financial indus-

try, we examined the process by which the technology was chosen, and how it is being deployed.

One product we studied was SmartRoles by Bridgestream. The software is "an enterprise-class application to capture, model, and update relationships between people, processes, projects, documents, locations, and business resources" [2]. The system seemed most focused on understanding "roles and organizational hierarchies." As a repository for relationship information and workflow interface for updating those relationships, one might expect a human resources team to be deploying it instead of a team of risk focused technologists. However, the application's sophisticated use of rules, which translate business policies into automated decision-making on entitlements, and the fact that this system will tie into existing technical entitlement solutions, clearly establish the need for technology professionals to evaluate and motivate its use.

The group of technologists who are behind the SmartRoles deployment first interacted with Bridgestream in the product's early stages. This recognition of the scope and importance of the problem in question was also mirrored by professionals at another partner firm, although at that time there was no product available to even come close to solving it [4]. Once establishing contact, the firm worked with Bridgestream over the course of more than a year to figure out what the product would need to include before being deployed in an operating environment. This delay, and corresponding willingness of both to negotiate features and interfaces over a surprisingly long period of time, reflects both Bridgestream's eagerness to have an important financial customer, and the firm's particularly open yet stringent philosophy regarding vendors, which we discussed previously in this document.

Once IT professionals decided to push a deployment of SmartRoles to the corporation, they had to establish sufficient buy-in among various business units. This emphasizes the cultural expectation that business groups are not forced to conform to IT initiatives; if the IT team felt that the application was good for the firm, they would have to make the case and convince non-technical employees, too. Because the IT team must get significant input and acceptance from the business group (who will essentially be the users of the application), the chance of deploying an unusable or imperfect product is even further reduced than if the IT team were evaluating it alone. This stands in stark contrast to the policies and customs of other companies and industries.

The technical team solicited buy-in from business groups using a well-rehearsed "road show," in which they collaborated with Bridgestream to demonstrate the product and the ways in which it would be useful to various groups within the company. Generally, it seems that IT personnel at the firm we were observing are unusually in touch with the needs and attitudes of their business counterparts, which makes the process of selling a product much easier. However, they felt that there was a distinct "terminology barrier" in discussing the roles rules at the heart of the product. In many cases, business and IT staffers have developed a common language to discuss existing notions and technology, but our collaborators found in this instance that the concepts necessary were both radically new, yet similar enough to existing business ideas to be confusing. For example, when discussing a role, does one mean a job title, the subtler part that an employee plays within the business group, or something else entirely? More generally, there seemed to be a difficulty in understanding the fundamental need

for such a product, the ways in which this firm's matrixed structure and demand an intricate and flexible system. We feel that this reflects a prescience of their technical staff: they are ahead of the curve in understanding the needs of their business counterparts, as well as in solving problems that other firms in their industry are still grasping to formulate.

The firm is now in the process of rolling out SmartRoles in a select number of systems within the firm. This initial deployment is critical to the success of the application, as it will determine whether or not other groups decide to adopt it. They have initially targeted systems that already have a strong notion and need for with entitlement, that will not require a lot of maintenance by the business teams but that will demonstrate the benefits and efficiency that can be gained.

5. Conclusions and Future Work

The results of our field research are indeed interesting. We have witnessed the challenges facing developers and managers in enabling appropriate information access in these data-driven firms. The changing organization environment, information accessibility, and regulatory environment all contribute to these challenges. In our ongoing research, we are developing models of both the organizational and system application structure to allow us to simulate the effectiveness of potential technical and access policy changes. For example, a model of an organization that allows the simulation of employee hiring, termination, promotion, and supervisory relationship changes could enable us to predict how auto-provisioning users with a certain role at a certain life-cycle event would affect the overall system. As part of this collaborative work with our partners in the financial industry, we plan to test our model with sanitized data and insight from the firm's organization.

Thus far, our collaboration with the field study firms has focused on role management. As we move forward in exploring the deployment of role technology, we also hope to examine the nature and usefulness of *rules* in entitlement systems. The firms we observed clearly recognize the benefits of rules in role-based systems, yet do not (nor does the research community) have a firm grasp on how to measure the effectiveness of a rule system, how to gauge the usability of a role-crafting interface, or how to integrate rule design with facilities for rule assignment, administration, and use monitoring. Furthermore, we have noted significant concern among all our financial research partners with "toxic combinations" of data, and how to recognize and prevent inappropriate access to them.

References

1. Anderson, Richard E. (1994). "Matrix Redux," *Business Horizons*, Nov.-Dec., 6-10.
2. Bridgestream, "Products," 2006. http://www.bridgestream.com/products.php.
3. Burns, Lawton R. and Douglas R. Wholey, "Adoption and Abandonment of Matrix Management Programs: Effects on Organizational Characteristics and Interorganizational Networks." *Academy of Management Journal*, Vol. 36, 1, 106-139.

4. Marc Donner, David Nochin, Dennis Shasha, and Wendy Walasek, "Algorithms and Experience in Increasing the Intelligibility and Hygiene of Access Control in Large Organizations," *Proceedings of the IFIP TC11/ WG11.3 Fourteenth Annual Working Conference on Database Security*. Kluwer, 2001.

5. Dynes, Scott, Eva Andrijcic, and M. Eric Johnson (2006), "Costs to the U.S. Economy of Information Infrastructure Failures: Estimates from Field Studies and Economic Data," *Proceedings of the Fifth Workshop on the Economics of Information Security*, Cambridge University, June.

6. Johnson, M. E., "A Broader Context for Information Security," *Financial Times*, September 16, 2005, 4.

7. Johnson, M. Eric and Eric Goetz (2007), "Embedding Information Security Risk Management into the Extended Enterprise," *IEEE Security and Privacy*, May-June, 16-24.

8. Kunreuther, H. (2002), "Risk Analysis and Risk Management in an Uncertain World," *Risk Analysis*, Vol. 22, No. 4, 655-664.

9. Kunreuther, H. and G. Heal (2003), "Interdependent Security," *The Journal of Risk and Uncertainty*, Vol. 26, No. 2, 231-249.

10. Li, N. and J. C. Mitchell, "RT: A Role-Based Trust-management Framework," *Proceedings of The Third DARPA Information Survivability Conference and Exposition (DISCEX III),* April 2003.

11. Ferraiolo, D. and R. Kuhn, "An Introduction to Robe-Based Access Control," NIST/ITL Bulletin, December 1995. http://csrc.nist.gov/rbac/NIST-ITL-RBAC-bulletin.html.

12. Sandhu, R., E.J. Coyne, H.L. Feinstein, and C.E. Youman, "Role-Based Access Control Models," *IEEE Computer*, 29 (2): 38-47, February 1996.

13. Smith, S.W., "Humans in the Loop: Human-Computer Interaction and Security," *IEEE Security and Privacy*, 1 (3): 75-79, May/June 2003.

14. Smith, S.W., C. Masone, and S. Sinclair, "Expressing Trust in Distributed Systems: the Mismatch Between Tools and Reality," *Forty-Second Annual Allerton Conference on Communication, Control, and Computing,* September 2004.